Cost- and Price Dynamics of Solar PV Modules

Stefan Reichelstein*
Graduate School of Business
Steyer-Taylor Center for Energy Policy and Finance
Stanford University

and

Anshuman Sahoo
Graduate School of Business
Steyer-Taylor Center for Energy Policy and Finance
Stanford University

April 2015

*Contact information: reichelstein@stanford.edu; asahoo@stanford.edu
We gratefully acknowledge financial support from the Steyer-Taylor Center for Energy Policy and Finance and a U.S. Department of Energy (DOE) grant administered through the Bay Area Photovoltaics Consortium (BAPVC). We also thank Aneesh Raghunandan for assisting us with data collection. All errors are our own.

Cost- and Price Dynamics of Solar PV Modules

Abstract: For several decades, the prices for solar photovoltaic (PV) modules have adhered closely to an 80% learning curve. Yet recent price declines have been even steeper. Analysts have questioned whether these price declines reflect underlying reductions in production cost or excessive additions to manufacturing capacity. For a sample of solar PV manufacturers, we estimate production costs based on financial accounting statements. We use these cost estimates as data inputs in a dynamic model of competition to obtain equilibrium prices, termed *Economically Sustainable Prices (ESP)*. We find that, in parts of 2012 and 2013, the ESPs significantly exceeded the observed average sales prices (ASP). At the same time, the observed dynamics of production costs point to reductions that are even faster than suggested by the 80% learning curve. Our estimates allow us to extrapolate a trajectory of future equilibrium prices to which ASPs should converge over time.

Keywords: Solar PV manufacturing, Industry dynamics, Economically Sustainable Price

JEL codes: D41, L11, L63, M21, Q42

1 Introduction

The solar photovoltaic (PV) industry has in recent years experienced rapid growth in the volume of output produced, sharp price declines for solar PV modules and a significant shift in the composition of module suppliers. To illustrate the growth dynamics, the 17 Gigawatts (GW) of new solar PV power capacity installed worldwide in 2010 was equal to the total *cumulative* installations of solar PV power over the previous four decades. Increasingly, this demand has been met by Chinese firms. These companies added significant production capacity while many solar manufacturers in the U.S. and Europe exited the industry.[1]

Swanson (2011) summarizes the price history from 1979 to 2010 and documents that this price trajectory conforms closely to the pattern of an 80% learning curve (Figure 1). Accordingly, prices have dropped by 20% with every doubling of cumulative output.[23] Particularly noteworthy is the 40% price drop in 2011 alone and the rebound in prices for late 2013.

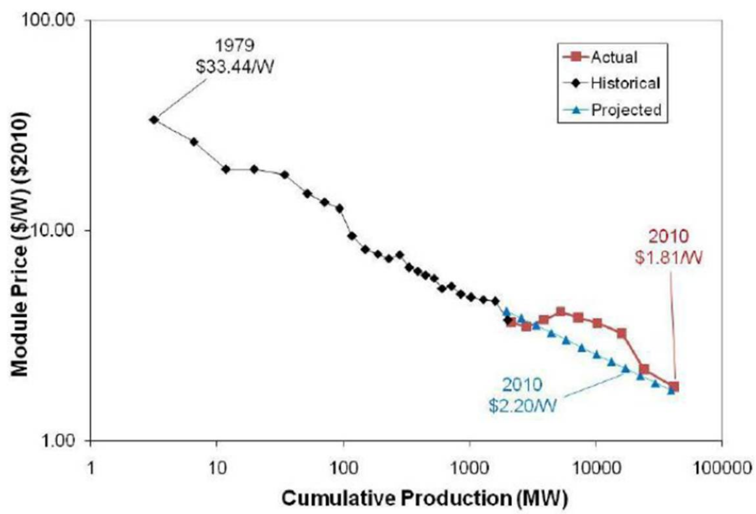

Figure 1: *Reproduction of plot from Swanson (2011)*

Between 2011 and 2013, the decline in average sales prices (ASP) for PV modules was

[1]In 2008, firms headquartered in the U.S. and Europe held 28% of module manufacturing capacity, while firms headquartered in China held 46%. By 2013, these figures had moved to 18% and 56%, respectively (Lux Research, 2013).

[2]Timilsina, Kurdgelashvili, and Narbel (2011) regress average sales prices (ASPs) on cumulative production volume over the period 1979-2010 and obtain a learning rate of 81%.

[3]Figures 1 and 2 show that for the years 2008-2009, ASPs were distinctly above the trend line suggested by the 80% learning curve that characterized the industry from 1979 to 2010. Most industry observers attribute this discrepancy to an acute polysilicon shortage which temporarily increased the raw material cost of silicon wafers.

even steeper than that predicted by the historical 80% learning curve, which we compare to ASP observations in Figure 2. Industry analysts have pointed out that the steep price declines in recent years may reflect in large part that the additions to industry-wide manufacturing capacity were excessive. If the latter explanation is valid, price forecasts based on mechanical extrapolations of recent price data would be misleading for manufacturers, solar power developers and policy makers. For example, the SunShot goals set by the U.S. Department of Energy (DOE) envision that a module price of \$0.50/W would allow solar power to be widely cost competitive. The anticipated timing with which this milestone will be reached may influence the duration of policies such as the 30% investment tax credit currently available in the U.S. to investors in solar power generation facilities.

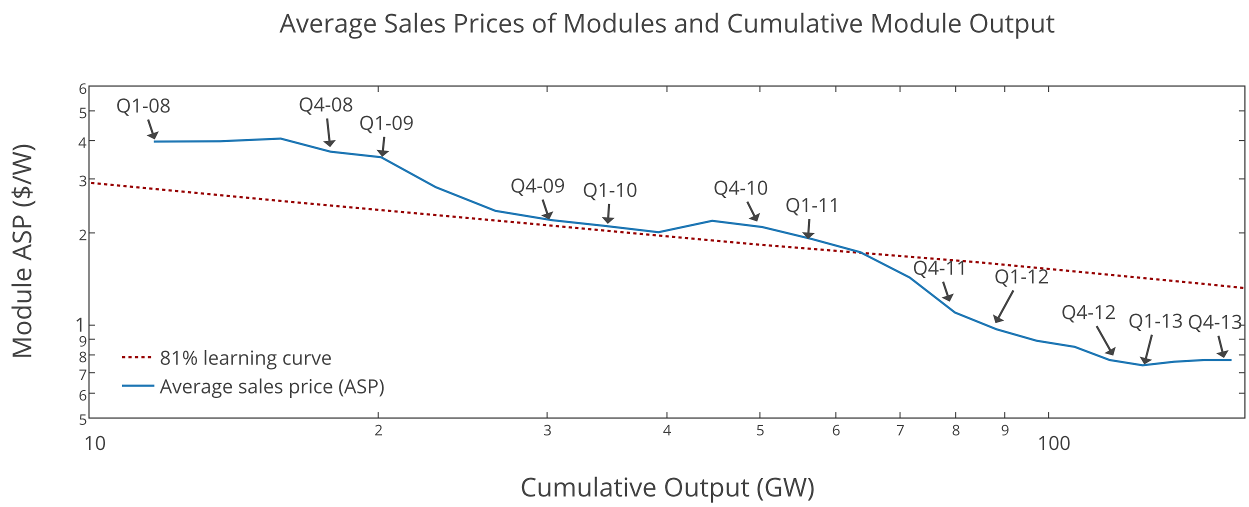

Figure 2: *Predicted and observed ASPs, 2008 – 2013. All prices are in 2013 U.S. dollars.*

Our paper makes several contributions to the debate on solar PV module prices. We estimate the production costs of module manufacturers in order to back out prices that would have prevailed between 2008 and 2013, assuming the industry had been in a competitive equilibrium. To do so, we develop a method for deriving production costs from financial statements. The resulting estimated equilibrium prices can be compared to the actually observed ASPs to test whether the industry was out of equilibrium at different points in time. Finally, our analysis of the production costs incurred during the years 2008-2013 can be used to extrapolate a trajectory of future production costs and corresponding equilibrium prices. This trajectory represents our benchmark of the industry fundamentals and can be interpreted as a trend-line to which actual prices should converge over time.

Our central cost measure is the long-run marginal cost of manufacturing and delivering one unit of output. Following industry convention, our unit of output is the number of Watts (W) of solar power in a module. The long-run marginal cost comprises capacity related costs in connection with machinery and equipment, current manufacturing costs for materials, labor and overhead as well as periodic costs related to selling and administrative expenses. Our formulation allows for both the cost of capacity and periodic operating costs to decrease over time due to exogenous technological progress and learning-by-doing effects.

We formulate a dynamic model of a competitive industry in which firms make a sequence of overlapping capacity investments and then choose their subsequent output levels in a competitive fashion, that is, taking market prices as given. While the long-run marginal cost contains components that are sunk in the short-run, the expected market prices will in equilibrium nonetheless be equal to the long-run marginal cost, because firms are capacity constrained. Furthermore, firms will earn zero economic profits on their capacity investments if the expected market prices in future periods are equal to the long-run marginal cost in those future periods. Accordingly, we refer to the long-run marginal cost at a particular point in time as the *Economically Sustainable Price* (ESP).

To address the question of whether the dramatic price declines for solar modules in recent years can be attributed in significant part to underlying cost reductions, we compare the average sales prices (ASP) for solar modules with the ESPs for 24 quarterly observations between 2008 and 2013. Our calculations focus on ten major module manufacturers with a combined market share of 35%.[4] All of these firms are listed on U.S. stock exchanges; therefore their financial statements have been prepared in accordance with U.S. accounting principles (GAAP). We infer the production costs for these firms from quarterly financial statements, including income- and cash flow statements as well as balance sheet data. In addition, we obtain data on manufacturing capacity and shipments from analyst reports and industry associations.

Our findings show a relatively close match between average sales prices and economically sustainable prices for the years 2008 – 2010. While our calculations reveal a steady and

[4]The solar PV industry satisfies the criteria of a competitive industry insofar as a large number of firms in the industry supply a relatively homogeneous product. To note, the median capacity market share of firms in this industry was 1% in 2012. Pillai and McLaughlin (2013) observe that firms are differentiated in terms of the efficiency of their modules, though for the firms in our sample the efficiency adjusted selling prices are quite similar.

significant decrease in costs and economically sustainable prices for the entire observation period, we also conclude that the dramatic decline in the observed ASPs for the years 2011-2013 is inconsistent with the industry having been in equilibrium during all periods of those years.[5] In other words, for those years the drop in ASPs should in large part be attributed to excessive additions to manufacturing capacity rather than to cost reductions.[6] For example, our estimates suggest an ESP of \$1.19/W at the end of 2012, but the ASP was about \$0.80/W.[7]

Despite our conclusion that observed prices were substantially below the ESP levels for at least some of the years, we also find that over our sample period the economically sustainable prices (the long-run marginal cost) declined even faster than suggested by the 80% learning curve. In particular, for production costs comprising materials, labor, and manufacturing overhead (excluding depreciation), we estimate a 74% learning curve. At the same time, we find that capacity related costs for machinery and equipment have fallen at a rate that also outperforms the 80% learning curve benchmark.

Our estimates of different learning parameters allow us to extrapolate the trajectory of future ESPs as a function of time and future production volumes. Under the conservative assumption that the industry will continue to produce and install about 40 GW annually in the coming years, we obtain a fundamental trend-line to which we expect ASPs to converge. Projecting forward to 2017 and 2020, we estimate ESPs of about \$0.70/W and \$0.61/W, respectively. While these values would represent a first-order effect on lowering the cost of solar power, they are significantly above the \$0.50/W price target of the DOE Sunshot Initiative.

The methods we employ in this paper are applicable beyond the solar PV industry. The identification of the ESP as the cost-based price which represents a competitive equilibrium price relates our work to that by Spence (1981) and Dick (1991). Unlike these earlier studies, we consider the role of new capacity acquisitions, the cost of which is expected to decline over

[5]Our conclusion is corroborated by the sharply negative earnings and declining share prices firms in our sample experienced during those two years.

[6]It is conceivable that these capacity additions reflected expectations about an expansion in demand that ultimately did not materialize.

[7]Our analysis allows us to conclude that, even in 2012, the observed ASPs covered those parts of the manufacturing costs that are typically considered "avoidable" in the short-run. This observation speaks to antitrust and dumping disputes in which suppliers are frequently held to a pricing standard that requires prices to cover all production costs that are avoidable in the short-run.

time. Similar models of overlapping capacity investments have been considered in Arrow (1964), Rogerson (2008), Rajan and Reichelstein (2009), Rogerson (2011) and Nezlobin (2012).

The cost inference method we employ to estimate the ESP provides an alternative to so-called "bottom-up" cost models, such as those in Powell et al. (2012), Powell et al. (2013), Goodrich et al. (2013a), and Goodrich et al. (2013b). These studies estimate costs by aggregating input requirements and prices as reported by various industry sources. Our approach based on reported accounting information can be viewed as a validation of the bottom-up cost models. In addition, our method yields equilibrium price predictions that account for anticipated future reductions in manufacturing costs. The work by Pillai and McLaughlin (2013) is closer in spirit to our analysis as it also applies firm-level accounting data to parametrize a model of competition in the solar manufacturing industry. However, as explained in more detail in Section 3 below, Pillai and McLaughlin rely on a measure of decision-relevant costs that differs fundamentally from the one in our paper.

The remainder of the paper is organized as follows. Section 2 formulates a dynamic model of a competitive industry with falling production costs. This framework allows us to identify Economically Sustainable Prices (ESP) in terms of production costs. Section 3 describes our data and inferential procedure. We then compare our ESP estimates to observed ASPs to test when the solar PV module industry was in equilibrium. Section 4 presents our econometric estimates of recent learning effects in manufacturing costs and applies these estimates to extrapolate a trajectory of future ESPs. Section 5 concludes. The appendix sections present proofs, data sources and adjustments used in our inference procedure, and robustness checks.

2 A Model of Economically Sustainable Prices

2.1 Base Model

The model framework we develop in this section allows us to identify economically sustainable prices in terms of production costs. We consider a dynamic model of an industry composed of a large number of suppliers who behave competitively. A key feature of the model is that firms are capacity constrained in the short-run. Each firm's output supplied to the market in a particular period is limited to the overall capacity that the firm has installed in previous periods. Production capacity available at any given point in time therefore reflects the cumulative effect of past investments, as in Arrow (1964), Rogerson (2008), Rajan and Reichelstein (2009), Dutta and Reichelstein (2010) and Rogerson (2011).

In the base version of the model, firms can accurately predict future demand. Let $P_t^o(Q_t)$ denote the aggregate willingness-to-pay (inverse demand) curve at time t, where Q_t denotes the aggregate quantity supplied at date t. Market demand is assumed to be decreasing in price and, in addition, we postulate that demand is expanding over time in the sense that:

$$P_{t+1}^o(Q) \geq P_t^o(Q), \tag{1}$$

for all $t \geq 1$ and all Q. The significance of this condition is that if firms make investments sufficient to meet demand in the short-run, they will not find themselves with excess capacity in future periods.[8] Given the growth in solar energy deployments, Condition 1 appears plausible in the context of solar PV modules.[9]

In order to break-even on their capacity investments, firms must realize a stream of revenues that covers periodic operating costs in addition to investment expenditures. The concept of an economically sustainable price is cost-based and comprises capacity related costs, periodic operating costs, and costs related to income tax payments. At the initial date 0, the industry is assumed to have a certain stock of capacity in place. To acquire one unit of manufacturing capacity, i.e., the capacity to produce one Watt (W) of solar PV modules per year, firms must incur an investment expenditure of v at the initial date 0. We allow for technological progress resulting in lower capacity acquisition costs over time. For reasons of tractability, though, we confine attention to a single "technological progress parameter", η, leading to a pattern of geometric declines such that $\eta^t \cdot v$ denotes the acquisition cost for one unit of capacity at time t, with $\eta \leq 1$.[10] Accordingly, investment decisions and the subsequent level of aggregate capacity in the market are conditional on firms' expectation of future decreases in capacity costs.

Investments in capacity represent a joint cost insofar as one unit of capacity acquired at time t will allow the firm to produce one unit of output in each of the next T periods.[11]

[8] Rogerson (2008) and Rogerson (2011) refer to (1) as the No-Excess Capacity (NEC) condition.

[9] Indeed, cumulative installed capacity has increased monotonically from 1.2GW in 2005 to 39GW in 2013 (Bloomberg New Energy Finance, 2014).

[10] Decreases in capacity cost as a function of time can be attributed to improvements in manufacturing equipment and improved efficiency for the conversion of modules.

[11] For simplicity, we adopt the assumption that productive capacity remains constant over the useful life of a facility. In the regulation literature, this productivity pattern is frequently referred to as the "one-hoss

To identify equilibrium prices in terms of costs, it will be useful to introduce the marginal cost of one unit of capacity *made available for one period of time.* As shown by Arrow (1964) and Rogerson (2008), this effectively amounts to "levelizing" the initial investment expenditure. To that end, let $\gamma = \frac{1}{1+r}$ denote the applicable discount factor. Practical capacity available in any period may only be a fraction of the theoretical capacity, and we denote the corresponding capacity factor by CF. The marginal cost of one unit of capacity in period t then becomes:

$$c_t = \frac{\eta^t \cdot v}{CF \cdot \sum_{\tau=1}^{T} (\gamma \cdot \eta)^\tau}. \tag{2}$$

An intuitive way to verify this claim is to assume that firms in the industry can rent capacity services on a periodic basis. Assuming this rental market is competitive and capacity providers have the same cost of capital, it is readily verified that the capacity provider who invests in one unit of capacity at time t and then rents out that capacity in each of the next T periods for a price of $c_{t+\tau}$ would exactly break even on his initial investment of $\eta^t \cdot v$. Accordingly, Arrow (1964) refers to c_t as the *user cost of capacity.*

In any given period, firms are assumed to incur fixed operating costs, e.g., maintenance, rent and insurance, in proportion to their capacity. Like past investment expenditures, these costs are assumed to be effectively "sunk" after date t because they are incurred regardless of capacity utilization. Formally, let f_t represent the fixed operating cost per unit of capacity available at time t, with $f_{t+1} \leq f_t$ for all $t \geq 1$. Finally, production of one unit of output entails a constant unit variable cost, w_t, which again is assumed to be weakly decreasing over time, that is, $w_{t+1} \leq w_t$ for all $t \geq 1$. In contrast to the fixed operating costs, variable costs are avoidable in the short-run if the firm decides not to utilize its available capacity.

Corporate income taxes affect the cost of production through depreciation tax shields and debt tax shields, as both interest payments on debt and depreciation charges reduce the firm's taxable income. Following the standard approach in corporate finance, we ignore the debt related tax shield provided the applicable discount rate, r, is interpreted as a weighted average cost of capital. The depreciation tax shield is determined by both the effective corporate income tax rate and the allowable depreciation schedule for the facility. These variables are represented as:

shay" model; see, for instance, Rogerson (2008) and Nezlobin, Rajan, and Reichelstein (2012).

- α : the effective corporate income tax rate (in %),
- d_t : the allowable tax depreciation charge in year t, $1 \leq t \leq T$, as a percent of the initial asset value.

The assumed useful life of an asset for tax purposes is usually shorter than the asset's actual economic useful life, which we denote by T in our model. Accordingly, we set $d_t = 0$ for those periods that exceed the useful life of the asset for tax purposes. As shown below, the impact of income taxes on the long-run marginal cost can be summarized by a *tax factor* which amounts to a "mark-up" on the unit cost of capacity, c_t.

$$\Delta = \frac{1 - \alpha \cdot \sum_{t=1}^{T} d_t \cdot \gamma^t}{1 - \alpha}. \tag{3}$$

The tax factor Δ exceeds 1 but is bounded above by $\frac{1}{1-\alpha}$.[12] It is readily verified that Δ is increasing and convex in the tax rate α. Holding α constant, a more accelerated tax depreciation schedule tends to lower Δ closer to 1. In particular, Δ would be equal to 1 if the tax code were to allow for full expensing of the investment immediately. We are now in a position to introduce the measure of long-run marginal cost:

$$ESP_t = w_t + f_t + c_t \cdot \Delta. \tag{4}$$

The label ESP_t on the left-hand side of (4) anticipates Finding 1 below, which shows that firms will break-even (achieve zero economic profits) over time if product prices are equal to ESP_t in each period. Given our assumption of a competitive fringe of suppliers, the investment and capacity levels of individual firms remain indeterminate. Denoting the aggregate industry-wide investment levels by I_t, the "one-hoss shay" assumption that productive assets have undiminished productivity for T periods implies that the aggregate capacity at date t is given by:

$$K_t = I_{t-T} + I_{t-T+1} + \ldots + I_{t-1} \tag{5}$$

Equation (5) holds only for $t > T$. If $t \leq T$, then $K_t = I_0 + I_1 + ... + I_{t-1}$.

[12] To calibrate the magnitude of this factor, for a corporate income tax rate of 35%, and a tax depreciation schedule corresponding to a 150% declining balance rule over 20 years, the tax factor would approximately amount to $\Delta = 1.3$.

Firms choose their actual output in a manner that is consistent with competitive supply behavior. Since capacity related costs and fixed operating costs are sunk in any given period, firms will exhaust their entire capacity only if the market price covers at least the short-run marginal cost w_t. Conversely, firms would rather idle part of their capacity with the consequence that the market price will not drop below w_t. Given an aggregate capacity level, K_t, in period t, the resulting market price is therefore given by:

$$p_t(K_t, w_t) = \max\{w_t, P_t^o(K_t)\},$$

while the aggregate output level, $Q_t(K_t, w_t)$ satisfies $P_t^o(Q_t(K_t, w_t)) = p_t(K_t, w_t)$. We refer to the resulting output and price levels as *competitive supply behavior.*

Definition 1 $\{K_t^*\}_{t=1}^{\infty}$ *is an equilibrium capacity trajectory if, given competitive supply behavior, the net present value of capacity investments at each point in time is zero.*

Finding 1 *The trajectory given by:*

$$ESP_t = P_t^o(K_t^*) \tag{6}$$

is an equilibrium capacity trajectory.[13]

The equilibrium price characterization in Finding 1 reinforces the interpretation of the ESP_t as the long-run marginal cost of one unit of output.[14] With additional assumptions, the capacity trajectory identified in Finding 1 is also the unique equilibrium capacity trajectory. This is readily seen if one assumes that capacity investments are reversible or, instead, that capacity can be obtained on a rental basis for one period at a time, with all rental capacity providers obtaining zero economic profits. Competition would then force the market price

[13] A formal proof of Finding 1 is presented in Appendix A. An implicit assumption here is that the aggregate capacity in place at the initial date 0 does not amount to excess capacity. Formally, we require $P_t^o(K_0) > ESP_1$.

[14] Our assumed competitive structure implies that no firm has a material impact on the probability of entry or exit by other firms. Thus, any mismatches between the predicted equilibrium prices and actual market prices stem from a market out of equilibrium rather than one with strategic predatory pricing as considered by, e.g., Besanko, Doraszelski, and Kryukov (2014). We emphasize that our model does not imply that firms are behaving myopically; rather, firms are forward looking and behave with the expectation of future cost reductions.

for the product in question to be equal to ESP_t in each period.[15]

In our model formulation, the unit variable costs, w_t, and the unit fixed cost, f_t, may decline over time, with the rate of decline taken as exogenous. We note that under conditions of atomistic competition, or that in which no firm can impact the prevailing market price through its own supply decision, the result in Finding 1 also extends to situations where the unit costs decline as a function of the *cumulative volume* of past output levels. One possible formulation is for $w_t = \beta(\sum Q_t) \cdot w$, where $\beta(\cdot) < 1$ is decreasing in its argument and $\sum Q_t \equiv \sum_{\tau \leq t} Q_\tau$.

To conclude this subsection, we note that our ESP concept is conceptually similar to the Levelized Cost of Electricity (LCOE) that is widely quoted in connection with the economics of different electricity generation platforms; see, for instance, Lazard (2009), Borenstein (2012), and Reichelstein and Yorston (2013). The LCOE yields a constant break-even price per kilowatt-hour that investors in a particular energy facility would need to receive *on average* in order to cover all costs and receive an adequate return on their initial investment. In contrast to our framework here, LCOE calculations are generally focused on the life-cycle of a single facility rather than on an infinite horizon setting with overlapping capacity investments which, per unit of capacity acquired, may become less expensive over time.

2.2 Demand Uncertainty

With one additional assumption, the characterization of equilibrium in Finding 1 can be extended to environments with price uncertainty. Suppose that given the aggregate supply quantity Q_t at date t, the price in period t is given by:

$$P_t(\epsilon_t, Q_t) = \epsilon_t \cdot P_t^o(Q_t),$$

where $\tilde{\epsilon}_t$ reflects volatility in the level of demand and is a random variable with mean 1. The support of $\tilde{\epsilon}_t$ is $[\underline{\epsilon}_t, \bar{\epsilon}_t]$, with $0 < \underline{\epsilon} < 1$. The noise terms $\{\tilde{\epsilon}_t\}_{t=1}^{\infty}$ are assumed to be serially uncorrelated, such that each $\tilde{\epsilon}_t$ is observed by all market participants at the beginning of period t. Competitive supply behavior then requires that:

[15] As our development makes clear, our estimates of future module prices are entirely cost based. This approach circumvents the challenges of endogeneity that would arise upon modeling the observed ASPs as a function of demand and installed capacity.

$$p_t(\epsilon_t, w_t, K_t) = \begin{cases} \epsilon_t \cdot P_t^o(K_t) & \text{if } \epsilon_t \geq \epsilon_t(K_t, w_t) \\ w_t & \text{if } \epsilon_t < \epsilon_t(K_t, w_t) \end{cases}$$

where the threshold level of demand volatility is given by:

$$\epsilon_t(K_t, w_t) = \begin{cases} \bar{\epsilon}_t & \text{if } \bar{\epsilon}_t \cdot P_t^o(K_t) \leq w_t \\ \frac{w_t}{P_t^o(K_t)} & \text{if } \bar{\epsilon}_t \cdot P_t^o(K_t) > w_t > \underline{\epsilon}_t \cdot P_t^o(K_t) \\ \underline{\epsilon}_t & \text{if } \underline{\epsilon}_t \cdot P_t^o(K_t) \geq w_t \end{cases}$$

Given K_t and w_t, the expected market price in period t then becomes:

$$E\left[p_t(w_t, \tilde{\epsilon}_t, K_t)\right] \equiv \int_{\underline{\epsilon}}^{\epsilon(K_t, w_t)} w_t \cdot h_t(\epsilon_t)\, d\epsilon_t + \int_{\epsilon(K_t, w_t)}^{\bar{\epsilon}} \epsilon \cdot P_t^o(K_t) \cdot h_t(\epsilon_t)\, d\epsilon_t \tag{7}$$

With risk neutral firms, price volatility will not affect the capacity levels obtained in equilibrium provided firms anticipate that they will exhaust the available capacity even for unfavorable price shocks. To that end, we introduce a condition of *limited price volatility*:

$$\underline{\epsilon}_t \cdot ESP_t \geq w_t.$$

Holding the distributions $h_t(\cdot)$ of $\tilde{\epsilon}_t$ fixed, this condition will be satisfied if the short-run avoidable cost w_t constitutes a relatively small percentage of the long-run marginal cost, ESP_t. The implication of this condition is that even for unfavorable price fluctuations firms will still want to deploy their entire capacity.

Corollary to Finding 1 *If price volatility is limited, the trajectory identified in Finding 1 remains an equilibrium capacity trajectory. The expected market prices in equilibrium satisfy:*

$$ESP_t = E\left[p_t(w_t, \tilde{\epsilon}_t, K_t^*)\right]. \tag{8}$$

As will become clear from the cost analysis in the following sections, the limited price volatility condition appears to be descriptive in the context of the solar PV module industry. Our estimates suggest that the unit variable cost, w_t, accounts for less than 60% of the total product cost, as measured by ESP_t, for the observation period 2008-2013. In the context of our model this would amount to $\underline{\epsilon}_t \geq 0.6$. To be sure, there are indications that some firms idled part of their available capacity during that time period, but we attribute this observation to the industry having been out of equilibrium, at least in parts of 2012 and 2013, rather than to the occurrence of large price shocks.

Finally, we note that the analysis in **?** has examined competitive equilibrium prices for settings where the limited price volatility condition does not hold. Their model presumes a finite horizon setting with T periods. In equilibrium, all firms make a single investment at the initial date and then supply output in a sequentially optimal manner, given periodic fluctuations in market demand. The sequence of expected equilibrium market prices still corresponds to the long-run marginal cost. Yet, the aggregate capacity level will generally be higher than that emerging under limited price volatility if firms anticipate that with some probability the aggregate capacity of the industry will not be exhausted. The intuition for this result is that since market prices are bounded below by the short-run variable cost, significant price volatility essentially entails a call option in the firm's payoff structure.[16] In order for the zero-profit condition to hold, the aggregate capacity level must, ceteris paribus, therefore increase.

3 Estimating Economically Sustainable Prices

This section presents our method for estimating the long-run marginal cost of production, the ESP, from financial accounting data. Our estimates for the cost of solar module production are on a per Watt basis. We begin by summarizing the main steps in the value chain of crystalline silicon PV modules. The major production steps consist of polysilicon, ingots, wafers, cells, and modules. The industry consensus is that opportunities for continued cost reductions remain at each step. Polysilicon is primarily produced via the so-called Siemens process, and the main cost reduction opportunities include improvements in energy efficiency and an increase in the scale of the Siemens chemical vapor depositor reactor.

The polysilicon output is used to grow ingots. Two recent developments have potential to reduce production costs. The first is the production of larger ingots. The second is the use of quasicrystalline ingots, as these eliminate a downstream processing step in which active silicon material is discarded. Quasicrystalline ingot growth could deliver \$0.10/W savings, relative to the incumbent technique for monocrystalline growth (Lux Research, 2012b). Ingots are subsequently sliced into wafers. The bulk of silicon losses occur at this

[16] Baldenius, Nezlobin, and Vaysman (2015) study a model of managerial performance evaluation in which negative shocks may also induce firms to leave parts of their capacity idle in some periods.

step, and the magnitude of losses depends on the thickness of the wire saws used. While smaller diameter wafer saws would reduce material loss during the process, these saws would be weaker and more likely to break.

The most capital and process intensive step of the module production process is the transition from wafers to cells. During this step, the wafers are etched and doped with impurities to achieve a desired level of electrical conductivity, metallized to facilitate the transfer of charges, and treated with an anti-reflective coating (Lux Research, 2012b). Finally, cells are strung together, enclosed, and appended with a junction box to build a solar module. Since module assembly requires only one essential piece of equipment, this last step has traditionally been labor intensive. Automation continues to reduce labor requirements substantially (Lux Research, 2012b).

We infer manufacturing costs from financial data released by Yingli Green Energy, Trina Solar, Suntech Power, Canadian Solar, LDK Solar, Hanwha SolarOne, JA Solar, ReneSolar, Jinko Solar, and China Sunergy.[17] Almost 300 other firms supply the module market (Lux Research, 2012a), but we excluded manufacturers based on five criteria. In particular, we excluded those firms (i) with lower than 0.5% share of global capacity in 2012, (ii) without public financial data since 2010, (iii) privately held or embedded within large conglomerates, (iv) listed on exchanges outside of the U.S, and (v) manufacturing thin-film modules. The firms in our sample manufacture primarily in China, are at least partially integrated across the value chain, and have invested in capacity expansions over our study period.[18]

Our data span 24 quarters from Q1 2008 to Q4 2013, and we use a time-index, $t \in \{1, ..., 24\}$, to refer to quarters in our inference technique. Table 1 presents the variables of interest. For simplicity, we have dropped the index i referring to individual firms. Nonetheless, since we infer manufacturing costs for multiple firms during each quarter, our sample includes 214 cost observations.[19]

All of the variables in Table 1 are obtained directly from individual data sources, except for quarterly production levels and additions to finished goods inventory, that is, q_t and n_t. In addition to these two variables, we also need to infer the cost of goods manufactured

[17]We access financial data through the Bloomberg terminal system. Table 6 in Appendix B provides summary details about the firms. The U.S. based firm SunPower is not in our sample because its sizeable solar development business makes it difficult to infer manufacturing costs from reported financial information.

[18]By inferring costs from financial accounting data, our cost estimates reflect subsidy benefits that may have been available to Chinese manufacturers.

[19]Since financial data are not available for all firms in all quarters, our panel is unbalanced.

Variable	Description	Units
q_t	Production volume in period t	MW_p
s_t	Sales volume in period t	MW_p
n_t	Inventory volume in period t	MW_p
$Sales_t$	Sales of modules during period t	\$
$COGS_t$	Cost of goods sold during period t	\$
$COGM_t$	Cost of goods manufactured during period t	\$
Inv_t	Value of inventory held at end of period t	\$
D_t	Depreciation charge in period t	\$
GM_t	Gross margin during period t	\$
ASP_t	Average selling price in period t	\$/W_p
$CAPX_t$	Capital expenditures in period t	\$
RD_t	Research and development expenses in period t	\$
SGA_t	Sales, goods, and administrative expenses in period t	\$
K_t	Production capacity in period t	MW_p/year
I_t	Production capacity addition in period t	MW_p/year

Table 1: *Variables obtained from financial data or derived from our inference procedure.*

($COGM_t$). Following industry practice, all of the above Watt measures are stated in terms of distributed current, that is, they are a calculated on a DC basis.

Our model framework in Section 2 conceptualizes the long-run marginal cost in each period as the sum of capacity related costs and current operating costs. Since we estimate these different cost components from firms' financial statements, we emphasize the basic distinction between *manufacturing* (inventoriable) costs and *period* costs. As the label suggests, the former pertain to factory-related costs, including materials, labor, and manufacturing overhead inclusive of depreciation. We consider all factory-related costs other than capacity costs as avoidable. Inventoriable costs are reflected in the figures reported for cost of goods sold (COGS) as part of the gross margin.[20] In contrast, *period costs* cover those related to selling as well as general and administrative (SG&A) expenses, including advertising and R&D. Conceptually, we think of w_t and f_t from Section 2 as having two components each:

[20]The sale by some firms in our sample of products other than solar modules and their upstream components introduces some noise to our COGS measure. However, these other products account for a small share of the firms' revenues. To illustrate, between 2008 and 2012, 96.5% to 98.7% of Yingli's revenues were from sales of modules. In 2013, RneSola's annual report splits sales data only according to modules and wafers.

$w_t = w_t^+ + w_t^-$ and $f_t = f_t^+ + f_t^-$, with the "+" part referring to manufacturing (inventoriable) costs and the "-" part referring to period costs.

Sections 3.1 and 3.2 present our inference procedure for manufacturing and period costs, respectively. The first presents our method for identifying those components of the unit variable and the fixed operating costs that are incurred in the manufacturing process. In the latter section, we derive our estimators of the costs comprising period costs.

3.1 Manufacturing Costs

3.1.1 Core Manufacturing Costs

We use the label *core manufacturing costs* to refer to all manufacturing (inventoriable) costs other than depreciation. Firms' financial accounting data allow us to make inferences only about the aggregate core manufacturing costs in period t and not about their fixed and variable components:

$$m_{it} \equiv w_{it}^+ + f_{it}^+,$$

The key variable for gauging manufacturing cost is Cost of Goods Manufactured (COGM). It is calculated as the unit cost m_{it} times the quantity of modules produced in the current quarter plus current depreciation charges pertaining to equipment and facilities:

$$COGM_{it} = Core\ Manufacturing\ Costs + Depreciation \equiv m_{it} \cdot q_{it} + D_{it}. \tag{9}$$

For our estimation purposes, only the depreciation charge, D_{it}, is directly available.[21] To estimate m_{it} in (9), we rely on several identities which connect production volume, sales, and inventory to infer the remaining variables, that is, $COGM_{it}$ and q_{it}. The quantity of units (modules on a per Watt basis) produced in period t equals the number of units sold plus the difference in inventory between the current and the prior period. Thus:

$$q_{it} = n_{it} - n_{it-1} + s_{it}. \tag{10}$$

[21] Depreciation charges are frequently reported on the annual statement of cash flows. When quarterly depreciation figures are unavailable, we apportion annual depreciation charges equally across quarters. That approach is consistent with the use of straight-line depreciation for financial reporting purposes.

Units sold in quarter t are sourced from both current period production and the inventory left from the prior period. Assuming that firms employ average costing for inventory valuation purposes, the average unit cost of firm i in period t is given by:

$$ac_{it} = \frac{Inv_{it-1} + COGM_{it}}{n_{it-1} + q_{it}}. \tag{11}$$

Here, ac_{it} is effectively the average cost per module available for sale by firm i at time t, taking the arithmetic mean between the beginning balance and the current period addition in both the numerator and the denominator. The left-hand-side of (11) can be inferred immediately from observations of cost of goods sold (COGS) and module shipments since:

$$COGS_{it} = s_{it} \cdot ac_{it}. \tag{12}$$

Finally, we make use of an expression yielding the ending inventory balance:

$$Inv_{it} = ac_{it} \cdot (n_{it-1} + q_{it} - s_{it}). \tag{13}$$

Adding this ending inventory balance to $COGS$, we obtain:

$$Inv_{it} + COGS_{it} = ac_{it} \cdot (n_{it-1} + q_{it} - s_{it}) + s_{it} \cdot ac_{it} = ac_{it} \cdot (n_{it-1} + q_{it}). \tag{14}$$

The right-hand side of (14) then allows us to infer $n_{it-1} + q_{it}$, which in turn yields $COGM_{it}$ in (9) via the fundamental identity in (10). In order to initialize the sequence of quarterly module production volumes, we infer an initial inventory level, n_{i0} by equating n_{i0} to $\frac{Inv_{i0}}{ac_{i0}}$.[22] Our inventory measure comprises finished goods and work-in-progress, but not raw materials.[23]

Finally, since firm-wide COGS and inventory apply to all products sold by the firm, we derive module-equivalent shipment levels.[24] To do so, we modify shipment levels for

[22] We index one quarter in each firm's data series to $t = 0$. The initial period is Q4-07 for most firms. The exceptions are LDK Solar and Jinko Solar, for which the initial periods are Q1-09 and Q2-10, respectively.

[23] Where quarterly breakdowns of inventory are unavailable, we assume that the split of inventory into finished and work-in-progress goods during the first, second, and third quarters is similar to the annual split observed in the previous and current years. Each quarter's split is a weighted combination of these two data points. The Q1, Q2, and Q3 estimates weight the previous year's annual data by 75%, 50%, and 25%, respectively.

[24] These resultant shipment levels thus account for the production and sale of other goods across the solar module value chain.

intermediate products (i.e., wafers and cells) by multiplying them by the ratio of their average selling prices in a quarter to that for modules. In particular:

$$s_{it}^{ME} = \phi_{W_t} \cdot s_{it}^{W} + \phi_{C_t} \cdot s_{it}^{C} + \phi_{M_t} \cdot s_{it}^{M}, \tag{15}$$

where $\phi_{W_t} = \frac{ASP_{W_t}}{ASP_{M_t}}$, $\phi_{C_t} = \frac{ASP_{C_t}}{ASP_{M_t}}$, and $\phi_{M_t} = 1$. The Online Appendix lists the multipliers we use.[25] To illustrate the equivalency concept with an example, consider a firm that ships 200MW of modules and 100MW of wafers in a given quarter. If ϕ_{W_t} were equal to 0.38, we would record the firm's module equivalent shipment quantity as $200 + 100 \cdot 0.38$ or 238MW. Upon doing so, we can use firm-wide COGS and inventory figures to estimate m_{it}.

Pillai and McLaughlin (2013) also use financial accounting data from solar PV manufacturers to infer production costs. While these data provide very useful guidance about the competitive dynamics in this market, our inference procedure differs markedly from theirs. First, Pillai and McLaughlin (2013) use Cost of Goods Sold (COGS) as their measure of variable cost of production. Since COGS includes depreciation charges and fixed operating costs, it will significantly exceed variable production costs. Secondly, Pillai and McLaughlin (2013) use the quotient of revenue to shipments in order to derive average selling prices (ASPs). Our inference procedure refines these estimates by adjusting shipments of wafers, cells, and modules in order to derive module-equivalent shipment measures.

3.1.2 Capacity Costs

For the purposes of calculating the long-run marginal cost of modules, we anchor the estimate of the capacity related costs to the baseline initial investment expenditure, v, required to manufacture one additional unit of output over the next T years. This expenditure is then "levelized" as shown in (2) in Section 2 to obtain the marginal cost of one unit of capacity made available at time t. Furthermore, this levelized cost must take into account the technological progress parameter η, which posits that the marginal cost of one unit of capacity made available at time t decreases geometrically with time.

Our analysis splits capacity related costs into two buckets: manufacturing equipment (e) and facilities (f). In accordance with (2), we obtain:

[25] The Online Appendix is available at http://stanford.io/1ov1kdQ.

$$c_t = c_{et} + c_{ft} = \eta_e^t \cdot \frac{v_e}{CF \cdot \sum_{\tau=1}^{T} (\gamma \cdot \eta_e)^\tau} + \eta_f^t \cdot \frac{v_f}{CF \cdot \sum_{\tau=1}^{T} (\gamma \cdot \eta_f)^\tau}. \tag{16}$$

Ideally, we would use firm-level data on fixed assets, depreciation, capital expenditures, and total capacity available to construct a quarterly panel of capacity costs. However, the data available entail several complications. First, it is unclear whether investment expenditures were directed at capacity upgrades or capacity additions. Second, the proportion of expenditure applied to investments in facilities as opposed to equipment is ambiguous. Additional concerns relate to inter-temporal allocation issues. One of these is the need to specify the lag between capital expenditures and capacity additions coming online. Finally, it is not clear how to split annual capital expenditure data into quarterly observations.[26] For these reasons, we estimate η by relying on data from a well-established industry observer, GTM (2012). Since one of our empirical goals is to estimate the distribution of ESPs across manufacturers and quarters, we do not directly use the capacity cost estimates from GTM, as they are aggregated across firms and at an annual level.

We first turn to the estimation of facilities costs. Since these costs pertain primarily to the cost of buildings, we set $\eta_f = 1$. However, as the efficiency of solar cells increases, the same physical area of output contains a greater power capacity (in Watts) and therefore the capacity cost per Watt decreases. We adjust the capacity cost in each period to reflect increases in average module efficiency. Appendix B details these adjustments. The top line of Table 2 records the efficiency-adjusted and industry-wide estimate of v_f; these numbers vary over time solely because of efficiency improvements. Our estimate of v_f is based on a bottom-up estimate of facility capacity costs for a known efficiency from Powell et al. (2013) and efficiency levels observed by Fraunhofer (2012) (see Table 7 in Appendix B). Powell et al. assume that a plant with an annual capacity of 395MW of 13.6% efficiency modules entails capacity costs of approximately \$53M, implying a v_f estimate of \$0.066/W.

We turn to equipment cost data provided by GTM (2012) to estimate the parameter η_e. The industry-wide estimates regarding v_e for different baseline years are summarized in Table 2. While the table presents equipment capacity numbers for China, GTM also provides cost

[26] A related challenge is that we observe changes in *total* capacity available, without being able to distinguish between additions to and retirements of capacity. However, since the useful life of equipment is generally estimated to be between 7 and 10 years and since the firms in our sample had relatively little capacity before 2005, we do not believe that this issue introduces a significant source of bias.

Component	2009	2010	2011	2012	2013	2014	2015	2016
Facility	\$0.07	\$0.07	\$0.06	\$0.06	\$0.06			
Ingot	\$0.21	\$0.17	\$0.11	\$0.09	\$0.06	\$0.05	\$0.04	\$0.04
Wafer	\$0.24	\$0.20	\$0.16	\$0.13	\$0.08	\$0.07	\$0.06	\$0.06
Cell	\$0.45	\$0.30	\$0.20	\$0.16	\$0.10	\$0.08	\$0.08	\$0.07
Module	\$0.12	\$0.09	\$0.07	\$0.05	\$0.03	\$0.03	\$0.02	\$0.02
Total Equipment	\$1.02	\$0.76	\$0.54	\$0.43	\$0.27	\$0.23	\$0.20	\$0.19

Table 2: *All rows but the first provide estimates from GTM (2012) of crystalline silicon capital equipment cost per watt in China. The first row lists our facility capacity cost estimates. Since average efficiency levels for 2014 – 2016 are still unknown, these entries are left blank. The last row presents the total equipment-related costs, excluding the cost of capacity for the facility.*

data for the U.S.; these costs are generally higher. Indexing each year by t and setting 2009 as year 0, we estimate the technological progress parameter via a simple regression:

$$v_t = \eta^t \cdot v_0 + \xi_t,$$

where ξ_t is assumed to be a log-normally distributed error term with $E\left[ln(\xi_t) \mid t\right] = 0$. Then, $log\frac{v_t}{v_0} = t \cdot ln(\eta_e) + ln(\xi_t)$. This regression estimate yields $\eta_e = 0.79$ (with a standard error of 0.02) for China and $\eta_e = 0.80$ (with a standard error of 0.02) for the U.S.

Since one of our goals is to derive a distribution of firm-specific ESP estimates with which to statistically test for a market in long-run equilibrium, we use our η_e estimate to derive firm-specific capacity costs. To align our capacity- and manufacturing cost estimates, we reset the index t to begin in 2007. We estimate equipment capacity costs for firm i, v_{ie}, as the quotient of cumulative capital expenditure, $CAPX$, between 2008 and 2013 to the change in module manufacturing capacity over that period. The technological progress parameter implies that in each subsequent year after 2007, the same capital expenditure yields more capacity per dollar. This prompts us to use the following adjusted measure of capital expenditures:

$$v_{ie} = v_{ie0} = \frac{\sum\limits_{t=0}^{6} CAPX_{it} \cdot \eta_e^{-t}}{K_{i2014} - K_{i2008}} \tag{17}$$

The η_e^{-t} term in the numerator "scales-up" capital expenditures after 2007 since such investments yielded more capacity per dollar.[27] The denominator, K_{i2014} - K_{i2008}, describes the change in manufacturing capacity.

Equation (17) is potentially prone to two erroneous inferences. First, our expression assumes that all capital expenditures were used to expand *integrated* module manufacturing capacity (i.e., capacity to produce ingots, wafers, cells, and modules). However, firms could have expanded their capacity to produce only some of these components. Second, firms' financial statements do not specify the portion of their capital expenditures that were applied to facility improvements as opposed to investments in new production equipment.[28]

We address the first issue by using an integrated module-equivalent (ME) level of capacity, K^{ME}, that "marks down" capacity additions that did not include all components of module manufacturing by the ratio of the capacity costs for the components actually installed to that for all components. Appendix B details our derivation of integrated module-equivalent capacity levels. The second issue leads us to define a factor β_t that measures the share of equipment capacity costs in total capacity costs:

$$\beta_t = \frac{v_e \cdot \eta_e^t}{v_e \cdot \eta_e^t + v_f}$$

The cost data in Table 2 allow us to calculate β_t on an annual basis. Since GTM does not provide capacity cost data for years preceding 2009, we backcast v_e for 2007 and 2008 by using our estimate $\eta_e = 0.79$. These two adjustments lead to the following modification of (17):

$$v_{ie} = \frac{\sum_{t=0}^{6} \beta_t \cdot CAPX_{it} \cdot \eta_e^{-t}}{K_{i2014}^{ME} - K_{i2008}^{ME}}. \tag{18}$$

We obtain firm-specific facility capacity costs by substituting $1 - \beta_t$ for β_t in (18). In general, our sample includes capital expenditure data for each year between 2007 and 2013 and capacity levels from 2008 to 2014.[29] In accordance with our model framework in Section

[27]Our approach thus yields estimates of the 2007 equipment capacity costs among the firms in our sample. Note that (17) assumes that capital expenses yield productive capacity within one year.

[28]A potential third issue is that we observe changes in *net* capacity and must implicitly assume that no capacity has been taken offline. This is a reasonable assumption, given the ages of the firms in our sample. Though we may be unable to observe deletions of capacity, we observe a decrease in capacity from one quarter to the next for only one firm and time period. If significant retirements actually occurred, our measure of v_{ie} would be biased upwards.

[29]Table 3 notes exceptions. More specifically, 2009 serves as the base year for JKS capital expenditures

2, we finally levelize the capacity acquisition costs to obtain a measure of cost per unit of output:

$$c_{iet} = \eta_e^t \cdot \frac{v_{ie}}{CF \cdot \sum_{\tau=1}^{10} (\eta_e \cdot \gamma)^\tau}, \tag{19}$$

and

$$c_{ift} = \frac{v_{if} \cdot \frac{eff_{ref}}{eff_t}}{CF \cdot \sum_{\tau=1}^{30} \gamma^\tau}. \tag{20}$$

The ratio of efficiencies in the latter equation reflects our adjustment for improvements in module efficiency over time (see Appendix B).

Since equipment and facility assets have different useful lives we employ separate tax factors, Δ_e and Δ_f, respectively. To calculate these, we apply a tax rate of $\alpha = 15\%$ and a (weighted average) cost of capital of $r = .13$. We set the useful life for equipment at 10 years and that for facilities at 20 years. These depreciation horizons are the minimum time horizons for equipment and buildings per Chinese law (PWC, 2012). Finally, we assume straight-line depreciation for tax purposes. These specifications are consistent with those in Goodrich et al. (2013b). Taken together, these inputs imply $\Delta_e = 1.08$ and $\Delta_f = 1.11$.[30]

Table 3 presents our estimates of levelized capacity costs, c_f, c_{e2007}, and c_{e2013}. The bottom row adjusts these values to incorporate the tax factors, that is, $c_f \cdot \Delta_f$ and $c_{e2013} \cdot \Delta_e$. The penultimate row in Table 3 corresponds to weighted averages for v_e, v_f, c_e, and c_f. These weights of firm-specific measures are calculated in proportion to each firm's share of capacity added between 2007 and 2014, relative to the total additions in the sample. The summary statistics do not include data from LDK or SOL, since these firms had significant investments in polysilicon capacity that could bias our estimates.

We summarize our estimates regarding the change in capacity costs as:

Finding 2 *Our estimates of the 2013 facility and equipment capacity costs are \$0.01/W and \$0.10/W, respectively. We estimate the technological progress parameter for equipment*

and 2010 for capacity. For STP, 2011 serves as the end year for capital expenditures and 2012 for capacity. For LDK, 2012 serves as the end year for capital expenditures and 2013 for capacity.

[30]Since the facility lifetime exceeds that of equipment, Δ_f can be adjusted to reflect the depreciation tax shield remaining on the facility investment at the time of equipment obsolescence. We do not do so, since the capacity costs related to facilities are small.

Firm	v_f, \$M/MW	c_f, \$/W	$v_e(2007)$, \$M/MW	c_{e2007}, \$/W	c_{e2013}, \$/W	Notes
CSUN	0.07	0.01	1.12	0.53	0.07	
YGE	0.12	0.02	2.14	1.01	0.14	
TSL	0.05	0.01	1.00	0.47	0.07	
JKS	0.05	0.01	0.74	0.35	0.05	Uses 2009 as base year
CSIQ	0.07	0.01	1.32	0.62	0.09	
HSOL	0.11	0.02	1.96	0.92	0.13	
LDK	0.18	0.02	3.85	1.81	0.25	Uses 2013 as end year
JASO	0.08	0.01	1.40	0.66	0.09	
SOL	0.12	0.02	2.02	0.95	0.13	
STP	0.07	0.01	1.61	0.76	0.25	Uses 2012 as end year
Weighted Avg.	0.08	0.01	1.42	0.67	0.09	
With tax factor		0.01		0.72	0.10	

Table 3: *Estimated cost of capacity, facility and equipment for firms in our sample set. Weights in the weighted average estimate are based on the share of module-equivalent capacity additions from 2007 to 2014.*

capacity costs to be η_e = 0.79, implying a reduction in equipment capacity costs of 21% per year.

In concluding this subsection, it should be noted that our procedure for inferring capacity acquisition costs yields estimates that are consistent with the data provided by GTM. In particular, we estimate a weighted average $v_e(2012)$ of \$0.42/W (not shown in Table 3), while the same figure from GTM equals \$0.43/W (see Table 2).

3.2 Period Costs

Having estimated the components w^+ and f^+ of the unit variable and fixed operating costs, w and f, this section develops estimates for the period costs, that is, the w^- and f^- components in our definition of ESP. As with our estimation of m_{it} in Section 3.1.1, we cannot identify w^- and f^- separately since we rely on financial accounting data. Period costs are primarily comprised of research and development (R&D) expenses and sales, general, and administrative (SG&A) costs. We treat R&D costs as an unavoidable fixed cost that provide the manufacturer an "entrance ticket" to internalize industry-wide cost reductions.[31] Selling

[31] Our assessment of firm-level R&D as an entrance ticket partly reflects the existence of large national R&D programs in China that sponsor centralized basic and applied research relevant to the solar industry

expenses are an example of a variable period cost, while administrative and managerial costs are examples of fixed period costs.

We observe firms' R&D and SG&A expenses directly from their income statements. To put these expenses on a per Watt basis, we divide each cost observation by the module-equivalent Watts of solar products produced by the firm in the given quarter. Referring back to the expression for the firm's ESP in (4) where $ESP_{it} = w_{it} + f_{it} + c_{it} \cdot \Delta$, we thus obtain:

$$ESP_{it} = m_{it} + \frac{(R\&D)_{it}}{q_{it}} + \frac{(SG\&A)_{it}}{q_{it}} + c_{it} \cdot \Delta. \tag{21}$$

The straightforward inclusion of period costs in (21) reflects an implicit assumption that these cost components have not been subject to learning effects. Over the period 2008 – 2013, we find that firms had median R&D and SG&A costs ranging from \$0.01/W – \$0.05/W and \$0.09/W – \$0.20/W, respectively.

3.3 Implied Economically Sustainable Prices

The inference procedure described in Sections 3.1 and 3.2 yields estimates for all the manufacturing- and period cost components of the long-run marginal cost. We are now in a position to calculate the economically sustainable prices, ESP, for each firm and quarter in our sample. These values in turn imply an industry-wide figure, ESP_t, which we define as the weighted average of the firm-specific ESPs in that period:

$$ESP_t = \sum_i w_{it} \cdot ESP_{it} \tag{22}$$

The weights w_{it} in (22) are derived on the basis of firm i's share of module-equivalent shipments in quarter t.

Figure 3 compares our ESP estimates with observed and in-sample ASPs, and we note three things.[32] First, before Q1-11, a casual inspection suggests that ASPs and ESPs generally stayed within a narrow band of each other. Second, the slope of the ESP and ASP

(e.g., the 863 and 973 programs). In future work it would be desirable to examine the impact of firm-specific R&D expenditures on the pace of future cost reductions obtained by individual firms.

[32]See Appendix B for details on the ASP series. The plot excludes the ESP estimate for SOL in Q4-08, given its value (\$57.92/W) far above any other in our estimation set. Though we exclude the point in the plot, we use it in hypothesis testing in this section.

curves match well until Q1-11, when the ASP curve decreases more steeply than the ESP curve. Third, the ESP and ASP measures appear to diverge after Q1-11. This latter period is one in which the market appears to have been out of equilibrium.

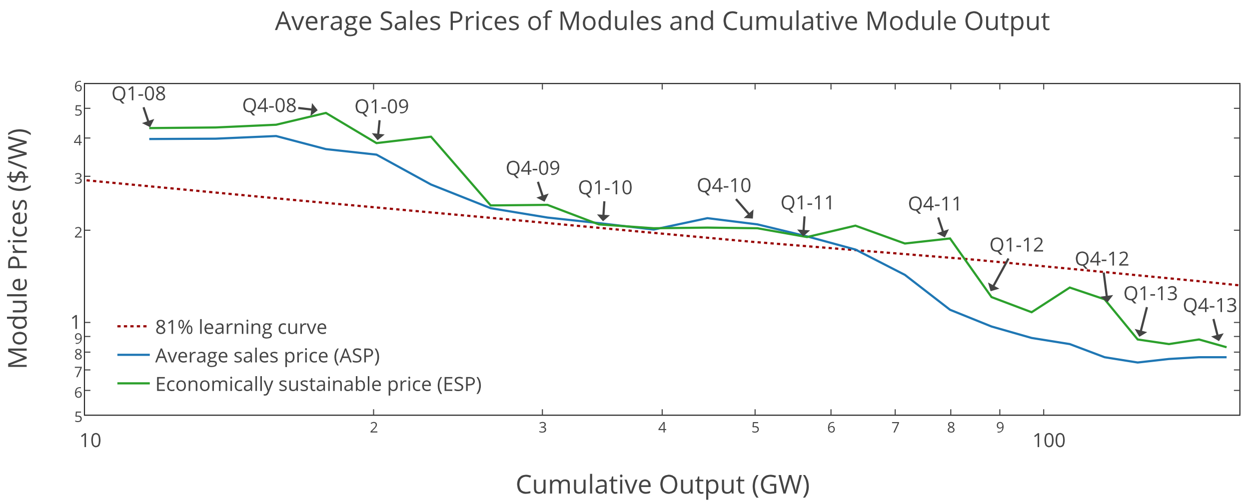

Figure 3: *ESPs generally track ASPs closely, though the two trajectories diverge after Q1-11. All prices are in 2013 U.S. dollars.*

Statistical inference allows us to make a formal claim regarding the pattern shown in Figure 3. To the extent that firm-specific ESPs and ASPs can be interpreted as draws from a distribution around the "true" market-wide ESP and ASP, the weighted standard deviation serves as a measure of the standard error around our ESP and ASP estimates and permits a statistical test of their equality.[33] For a given quarter, our procedure tests the null hypothesis that $ASP_t = ESP_t$. We perform two-tailed hypothesis tests with the alternative hypothesis being $ASP_t \neq ESP_t$.[34]

At the 10% significance level, we reject the equality of the ASP and ESP in Q1-12 (p =

[33]Note that we only use firm-specific ASPs, which are defined as the ratio of firm-specific revenues to module-equivalent shipments (see Appendix B). The mean in-sample market-wide ASP and ESP measures are the weighted averages of the firm-specific ASPs and ESPs, respectively, with weights defined by the share of shipments by firm i. We do not use price index data because we do not observe the distribution of prices around the index price.

[34]Per standard econometric texts (e.g., Greene (2003)), the non-overlap of a confidence interval for a given significance level of a random variable with zero is equivalent to a formal parametric statistical test at the same level. Thus, the test suggested here can be implemented by comparing confidence intervals around our mean ESP and ASP measures. However, the direct comparison of confidence intervals increases the chance of a type II error. We use the analytic derivation by Afshartous and Preston (2010) to calculate confidence intervals around the ESP and ASP measures and perform the equivalent of a t-test of the null hypothesis.

0.030), Q3-12 (p = .042), Q1-13 (p = .076), and Q2-13 (p = 0.093).[35] The Online Appendix details the weighted means, standard errors, and degrees of freedom used in our statistical tests. Our results support the claim that the market was out of equilibrium in 2012 and 2013. In addition, the statistical results suggest that, despite a tight polysilicon market in 2008, cost and price data from that year are consistent with a module market in equilibrium.[36] We summarize our inference regarding equilibrium in Finding 3:[37] While these results confirm the widely held belief that the industry experienced overcapacity in 2012 and 2013, our approach allows the analyst to quantify the extent to which price reductions stemmed from overcapacity as opposed to continued cost reductions. By developing a test for a market in long-run equilibrium, we also provide the analyst a means to identify quarters during which price data diverge from the fundamental cost trend.

Finding 3 *Our estimated ESP are statistically significantly different from the observed ASPs in Q1-12, Q3-12, Q1-13, and Q2-13. We thus infer that the solar PV module market was out of equilibrium during those periods.*

Though our data are consistent with the suggestion of overcapacity in the PV module market, we also observe a fall in ESPs from $1.82/W in Q1-11 to $0.82/W in Q4-13. This suggests that "true" cost reductions also explain part of the price decreases observed over this period. In Section 4 we therefore estimate the size of learning effects within the industry in order to derive a trajectory of future equilibrium module prices.

To conclude this section, we relate our measure of production costs to the so-called "Minimum Sustainable Price" (MSP) estimates in Powell et al. (2013) and Goodrich et al. (2013a,b). The MSP is similar to our ESP concept insofar as the Minimum Sustainable Price also seeks to identify a cost-based sales price that provides an adequate return to investors. In contrast to our top-down approach based on firm-level financial data, Goodrich et al. (2013a,b) rely on a bottom-up cost model in which individual cost components are assessed

[35]With a more stringent 5% significance level, we reject the null hypothesis of equality between the ASP and ESP for Q1-12 and Q3-12.

[36]A caveat is that our sample is smaller in the earlier period.

[37]Inference using the simple average of ASPs and ESPs broadly agrees with our results and implies a market out of long-run equilibrium in 2012 and 2013. At the 5% level, tests with the simple averages reject the equality of the ASP and ESP in Q4-09, Q3-11, Q1-12, Q2-12, Q3-12, Q1-13, Q2-13, and Q4-13. At the 10% level, the test additionally rejects the null hypothesis in Q4-11, Q4-12, and Q3-13.

in 2012 on the basis of various information sources available from industry observers. The MSP is then calculated as the derived manufacturing cost plus a profit mark-up.

We regard our top-down approach to the derivation of ESPs as complementary to the bottom-up cost models. The most important difference is that our formulation is inherently dynamic insofar as the long-run marginal cost (ESP) is assumed to change over time due to learning-by-doing and technological progress. These anticipated cost reductions apply both to core manufacturing costs and to capacity related costs and our notion of competition postulates that equilibrium prices incorporate these expected cost declines.

4 Forecasting Economically Sustainable Prices

A central motivation for the approach developed in this paper is to obtain a prediction model for future market prices that goes beyond a mere extrapolation of ASPs observed in the past. This section develops such a prediction model in the form a trajectory of future ESPs. The trajectory estimation requires projections both for capacity- and current manufacturing costs. Given our estimate of η_e from Section 3, the former is straightforward: for any time period, capacity costs are determined by the time elapsed since the period in which the baseline cost of capacity was calculated. We project future manufacturing costs by estimating a learning curve in Section 4.1. Section 4.2 combines our capacity cost decline parameter and manufacturing cost learning curve estimates to project ESPs through 2020.

4.1 Projecting Core Manufacturing Costs

As indicated in Section 2, our model framework allows for the possibility that operating costs decline with cumulative production volume, as posited in traditional learning curve studies. The results reported in this section represent the first attempt to estimate the decline in module manufacturing costs, specifically the core manufacturing costs, m_{it}. Our estimation also offers a useful update to existing learning curve estimates, as the latter use data from a period during which large changes occurred on both the demand- and supply sides. Through the early 1980s, the demand-side was characterized by a monopsony, with governmental programs purchasing the majority of solar modules (Swanson, 2011). On the supply side, the industry has shifted to a diversified set of suppliers, with a large fringe of small producers. In addition, solar PV manufacturers have overtaken semiconductor manufacturers in their collective demand for polysilicon. This implies first that reductions in upstream polysilicon production can be at least partially attributed to increased demand by solar PV manufacturers and second that, during a supply crunch during the years 2007 – 2009, solar PV manufacturers faced large incentives to improve their production processes.

4.1.1 Econometric Specifications

Equation (23) presents our base specification, Specification 1. Following our approach in Section 3.1.1, our dependent variable is the aggregate core manufacturing cost, $m_{it} = w_{it}^{+} + f_{it}^{+}$. We assume that m adheres to a constant elasticity learning curve.

$$m_{it} = m_{i1} \cdot (\frac{Q_t}{Q_1})^{-b} \cdot e^{\varepsilon_{it}} \tag{23}$$

In (23), Q_t is the in-sample cumulative production level in time t, b is the learning elasticity, and ε_{it} is an idiosyncratic error term, with $E\left[\varepsilon_{it} \mid Q_t\right] = 0 \ \forall \ i, t$.[38] The slope of the constant elasticity learning curve, S, is given by $S = 2^b$ (Lieberman, 1984). Given a projected production level at time t and an original manufacturing cost, m_{i1}, we can use S to forecast the manufacturing cost at that time.

Though most empirical studies find that static scale economies are small in magnitude when compared to learning effects, we follow their lead in explicitly controlling for changes in the scale of manufacturing facilities. Specification 2 introduces scale effects by assuming that manufacturing costs change exponentially with the scale of plants.[39] In particular, and as seen in (24), we introduce a term, $\Delta Scale_{it}$, equal to the difference between scale at time t and in Q1-08, and a corresponding coefficient, b_{scale}. $Scale_{it}$ is measured in MW/year and is defined as the average capacity per manufacturing site operated by firm i.

$$m_{it} = m_{i1} \cdot (\frac{Q_t}{Q_1})^{-b} \cdot e^{b_{scale} \Delta Scale_{it}} \cdot e^{\varepsilon_{it}} \tag{24}$$

Ideally, we would introduce time fixed effects or a trend in our specifications. However, as common in learning models, the correlation between time and cumulative output is high, and the inclusion of either time fixed effects or a temporal trend inflates standard errors to levels

[38]The in-sample production includes production by the firms in our sample set only.

[39]We test an alternative form of Specification 2 with scale effects structurally similar to the learning effects. This follows specifications used by Lieberman (1984) and Stobaugh and Townsend (1975) but yields a poorer fit to the observed ESP data than does the one reported.

precluding statistical inference.[40] Time fixed-effects can capture the impact of changes in input prices that may account for changes in manufacturing costs marginal to learning effects. Polysilicon is a major input in module manufacturing, and Figure 4 documents changes in polysilicon prices during our sample period. In principle, one could use polysilicon price data to control for changes in this input price, but the temporal lag between polysilicon procurement and utilization is not obvious.[41]

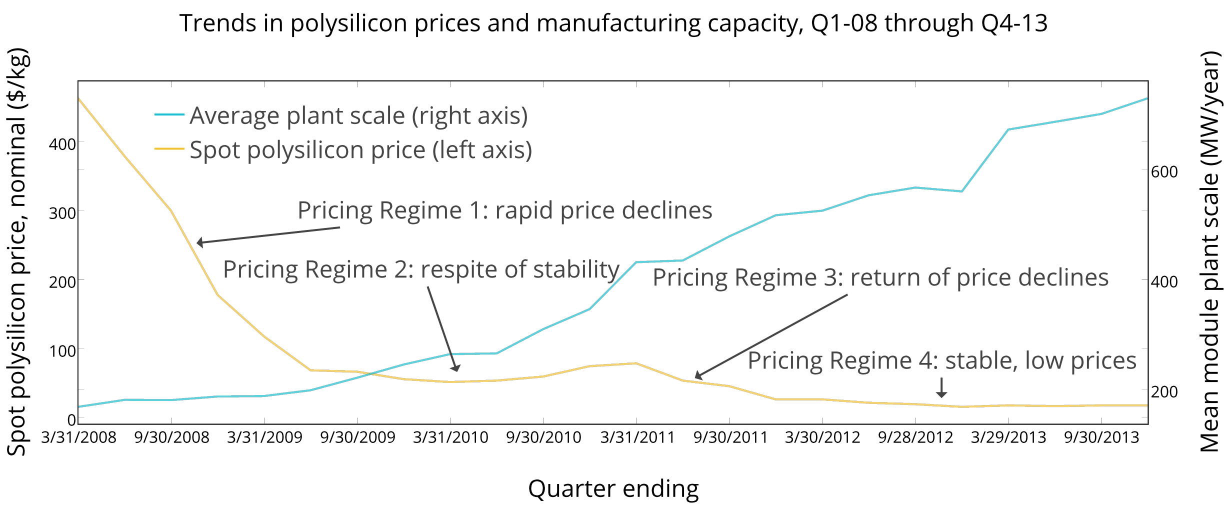

Figure 4: *As cumulative production increased between 2008 and 2013, the average size of module manufacturing facilities increased and the price of polysilicon decreased.*

Instead, we exploit in Specifications 3 and 4 the variation in the slope of polysilicon prices. Though Specification 3 is structurally the same as Specification 2, it uses data only from the time periods over which polysilicon prices remained relatively constant; these are labeled as Pricing Regimes 2 and 4 on Figure 4.[42] Specification 4 adds a Pricing Regime 4 dummy term to Specification 3 and yields a learning curve slope estimate that can be interpreted as an upper bound. By including this dummy term, we effectively 'remove'

[40]The correlation between time and cumulative production is 0.97. Cumulative production and scale have a smaller correlation of 0.65.

[41]Using a solar grade polysilicon price index from Bloomberg New Energy Finance (2014), we test models including a linear polysilicon price term from (1) the current time period, (2) one lagged period, (3) two lagged periods, and (4) a weighted average of two lagged time periods, where the weights are defined by the share of production in each quarter. While we do not present the resulting estimates, we find that the statistical significance of polysilicon prices is sensitive to the assumed lag structure.

[42]Pricing Regimes 1, 2, 3, and 4 include quarters Q1-08 through Q2-09, Q3-09 through Q4-10, Q1-11 through Q4-11, and Q1-12 through Q4-13, respectively. In our dataset, we include dummy terms for each regime that equal 1 for all records in the appropriate quarters.

all cost reductions between the end of Pricing Regime 2 and the start of Pricing Regime 4 from the estimate of the learning elasticity. If observed cost reductions were due only to polysilicon price decreases, we would expect the coefficient on cumulative output to be statistically indistinguishable from 0 in Specification 4.

While Specifications 3 and 4 compensate for our limited ability to include time fixed-effects, they do not substitute for temporal trends. The latter are important to the extent that the quality of output changed over time. Appendix C presents estimates in which we explicitly transform our cost data to a standard efficiency level. The only difference observed when estimating specifications with a standard efficiency level is that the learning curve slopes are slightly more gradual.

To use standard linear panel data methods, we state the model in its linear in logarithms form; for example, we express Specification 2 as in (25):

$$log(m_{it}) = log(m_{i1}) - b \cdot log(\frac{Q_t}{Q_1}) + b_{scale} \cdot \Delta Scale_{it} + \varepsilon_{it} \tag{25}$$

Across all specifications, we index observations from Q1-08 with $t = 1$.[43] Our cumulative output measure is based on the sum of estimated production across firms in our sample. We use firm-level plant scale data from Lux Research (2014), though we modify some records to reflect financial data and press reports. Across specifications, we assume the idiosyncratic error term has mean zero and is uncorrelated with the explanatory variables. Notably, we do not assume that $E[\varepsilon_{it}|m_{i1}] = 0$, and we accordingly use a fixed effect estimator.[44]

4.1.2 Details on Estimation and Inference

We make two sets of comments about our inference procedure that are relevant to the interpretation of our estimates. Relate to the data themselves, we assume that measurement errors introduced by our cost inference procedure are normally distributed with mean zero and uncorrelated with our explanatory variables; this implies they are captured by the idiosyncratic error terms. Since we use data from firms listed on U.S. exchanges, our estimates

[43]For two firms, JKS and LDK, our firm-specific production estimates do not cover Q1-08 since those firms did not release financial accounting data from these periods. To avoid dropping all observations in which those firms' data are unavailable, we use module production levels recorded in Lux Research (2014). For LDK, we use data from Lux Research (2014) for production observations between Q1-08 to Q1-09; for JKS, we use data from the same source for observations between Q1-08 and Q2-10.

[44]This allows for unobserved variation across firms that can explain variation in manufacturing costs.

are potentially subject to sample selection bias. Despite this, we avoid making strong assumptions required to estimate a parametric Heckman model. Our estimates of learning effects should be interpreted as conditional on public listing on a U.S. exchange. If our intuition that firms with the highest propensity to learn will be publicly listed is correct, the estimated conditional learning curve slope estimate would be weakly steeper than the unconditional version.

The second set of comments relates adjustments of the standard errors and inference procedure. To account for departures from homoskedasticity, auto-correlation within firms and cross-sectional dependence across them, we report standard errors suggested by Driscoll and Kraay (1998) and implemented by Hoechle (2007).[45] Finally, given the small size of our dataset, we correct the standard errors by scaling the asymptotic estimates by $\sqrt{\frac{N}{N-1} \cdot \frac{T-1}{T-k}}$, where N, T, and k are the number of firms, time periods, and coefficients, respectively. Moreover, our inference is based on a Student's t distribution with (N-1) degrees of freedom, rather than a standard normal distribution, to account for our small sample size.

4.1.3 Econometric estimations and interpretation

Table 4 presents our estimates.[46] Across all specifications, the coefficient on cumulative output is significant, while that on scale is not.[47] Taken together, Specifications 1 through 4 imply that significant cost decreases occurred between 2008 and 2013 and that these were not solely attributable to decreases in polysilicon prices.

Our preferred specification is Specification 3 because it accounts for possible unusual polysilicon pricing events in Pricing Regime 1 and controls for scale. Though Specification 4 does the same, we believe it is too conservative in excluding all cost reductions that

[45] Although the calculation of these standard errors relies on large sample asymptotics, the Driscoll-Kraay errors have better small-sample properties than common alternatives, such as cluster robust variance estimators (CRVE), when cross-sectional dependence exists (Hoechle, 2007).

[46] The magnitude of estimated coefficients and standard errors on firm scale is indeed smaller than 0.000.

[47] Our measure of cumulative output imputes production levels for Suntech Power (between Q2-12 and Q4-13) and LDK Solar (between Q1-13 and Q4-13). These firms were de-listed in Q2-12 and Q1-13, respectively, and we do not observe their subsequent production volumes. Our imputation assumes that these firms retain the 0.83% and 0.46% shares of global production, respectively, that we observe in the quarters for which we have financial accounting data. Without this imputation, our data would include a smaller increase in cumulative production than actually occurred, and we would estimate a coefficient on cumulative production that would be biased downwards. An estimation without imputed production from these two firms yields learning curve slopes that are 1% – 1.5% steeper than that reported in Table 4.

	Dependent variable: Log Manufacturing Cost/Watt			
Specification	**1**	**2**	**3**	**4**
Intercept	1.488***	1.440***	1.450***	0.767***
	(0.166)	(0.166)	(0.202)	(0.170)
Cumulative Production (b)	-0.427***	-0.390***	-0.433***	-0.193**
	(0.048)	(0.058)	(0.053)	(0.060)
Firm Scale (b_s)	–	-0.000	-0.000	0.000
		(0.000)	(0.000)	(0.000)
Dummy, PS Regime 4	–	–	–	-0.521***
				(0.117)
Learning Curve Slope (S)	74.4%	76.3%	74.1%	87.5%
	(15.6%)	(20.7%)	(16.9%)	(45.4%)
Adjusted R^2	0.7542	0.7593	0.7725	0.8157
N	213	213	125	125

Table 4: *Estimated coefficients for a constant elasticity learning curve. The intercept should be interpreted as the average of the logarithm of the Q1-08 core manufacturing cost across firms. Entries in parentheses are Driscoll-Kraay standard errors. Cumulative production is total output by firms in our sample.*
Key to statistical significance: ***: ≤ 0.001; **: ≤ 0.01; *: ≤ 0.05.

occurred while polysilicon prices declined in Pricing Regime 3, especially since over 90% of demand for polysilicon is from the solar market. Nonetheless, if one believes that the upstream polysilicon market is unlikely to achieve future cost reductions at the same rate characterizing it from Q2-09 through Q4-13, Specification 4 provides the appropriate rate of cost declines for core manufacturing costs. Importantly, since our preferred specification uses observations from quarters during which we observe evidence of off-equilibrium actions (i.e., overcapacity), we reiterate that our estimated parameters are based on cost, not price, observations.

Finding 4 *Controlling for plant scale and excluding periods with substantial polysilicon price declines, we estimate a 74% learning curve for core manufacturing costs over the period 2008-2013.*

Table 4 indicates that the propagation of error from our scale coefficient estimate to our learning curve estimate implies a standard error of approximately 17% on our learning curve slope. We account for this uncertainty by providing a 95% prediction interval around our projected ESPs in Section 4.2. Appendix C provides details of robustness checks on our inference about decreases in manufacturing costs.

4.2 ESP Projections

With our estimates of the technological progress parameter characterizing capacity cost dynamics and the learning curve slope applicable to core manufacturing costs, we are in a position to project future ESPs. Recall from Section 2 that we expect ASPs to converge to these ESPs as demand increases and both capacity- and core manufacturing costs decrease. The timing of convergence will depend on the trajectory of market demand. To illustrate the sensitivity of ESP projections to demand, we present ESP forecasts contingent on annual demand ranging from 40GW/year to 60GW/year. Our baseline forecast is based on a 40GW/year demand assumption, as it corresponds to observed demand in 2013.

To form our projections, we assume that a representative firm maintains the 35% share of global module production that the firms in our sample held in 2012. Accordingly, the increase in "within sample" cumulative production is $0.35 \cdot 40$GW, 50GW, or 60GW, depending on our demand assumption. Together with our estimated coefficient on cumulative production, we derive an expected core manufacturing cost for each year. Our projected capacity cost for the representative firm reflects a weighted average of firms' projected capacity costs for each of the years between 2014 and 2020, with weights determined by a given firm's share of 2013 shipments. Finally, we add R&D and SG&A costs that are equal to the 2013 shipment-weighted average of firms' median R&D and SG&A costs from Q1-08 to Q4-13.

Figure 5 depicts the forecast trajectory of ESPs through 2020. This curve entails a 27% reduction in production costs with every doubling in industry-wide output. This reflects both the 74% learning curve on core manufacturing costs and the 78.5% annual geometric decline in capacity costs, which occurs faster than does a doubling in output. To reflect the error present in our estimates, Table 5 presents 95% prediction intervals for each year under each of the three demand scenarios.

Our ESP projections prompt four comments. First, even under our most aggressive assumptions about future demand, our point estimates of future production costs imply that, if manufacturers were to earn zero economic profits, module prices would be unlikely to decrease to the 2020 SunShot target of \$0.50/W set by the U.S. Department of Energy.

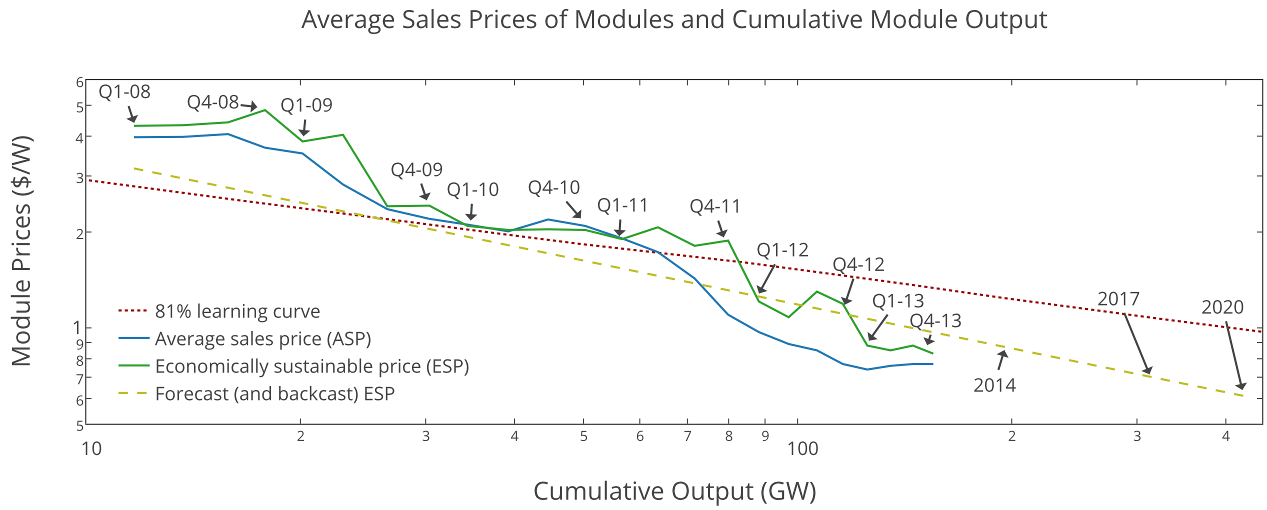

Figure 5: *Projected ESPs through 2020, assuming a constant yearly addition of 40GW. All prices are in 2013 U.S. dollars.*

Nonetheless, the \$0.50/W target is covered by our 95% prediction interval by 2016 or 2017 in all three demand scenarios. Second, the lowest ASP we observe in our price data is \$0.74/W in Q1-13. Under our base assumptions about future demand (i.e., demand additions of 40GW per year), we do not project ESPs at that level until 2016. This agrees with our expectation that solar module prices will not decrease monotonically but instead converge to the ESP trajectory. Once such convergence has occurred, new capacity may enter the market and introduce a new generation of module manufacturing technology. Third, our point estimates suggest that technological breakthroughs may be needed to push solar module prices toward SunShot goals. Finally, we note a need for a continued re-evaluation of the cost dynamics of solar modules; in an ideal case, the analyst would re-estimate learning parameters as new financial accounting data were released and use these updated estimates to update the expected ESP trajectory.

Our projected solar module prices are relevant to debates about policy support for solar energy. For example, the U.S., federal policy that allows investors in solar installations a 30% investment tax credit (ITC) is set to be reduced to 10% by the end of 2016. Comello and Reichelstein (2015) examine the likely impact of this change in policy on the competitiveness of solar power for various regions of the U.S. and for residential-, commercial- and utility-scale installations. Since future module prices are unlikely to decline as sharply as they did between 2008 and 2013, Comello and Reichelstein project that the planned step-down in the

	2014 ESP			2017 ESP			2020 ESP		
Demand	PE^-	**PE**	PE^+	PE^-	**PE**	PE^+	PE^-	**PE**	PE^+
40 GW	\$0.65	\$0.88	\$1.11	\$0.48	\$0.70	\$0.93	\$0.39	\$0.61	\$0.82
50 GW	\$0.64	\$0.87	\$1.10	\$0.46	\$0.68	\$0.90	\$0.37	\$0.58	\$0.79
60 GW	\$0.63	\$0.86	\$1.09	\$0.44	\$0.66	\$0.88	\$0.36	\$0.56	\$0.77

Table 5: *ESP projections, given assumptions about the yearly demand for solar PV modules. All figures are in 2013 dollars and are expressed per Watt. PE indicates "point estimate", "PE^+", the upper bound of the prediction interval and "PE^-", the lower bound of the prediction interval.*

federal ITC would leave solar power uncompetitive by early 2017 in nearly all segments and locations. The authors evaluate an alternative gradual reduction of the federal ITC between 2017 and 2024. The slower phase-down is calibrated to offset anticipated cost reductions in new solar installations and would envision an end to all federal tax support by 2025.

5 Conclusion

Prices for solar PV modules have dropped consistently over the past four decades and have historically adhered to an 80% price-based "learning curve." More dramatic price decreases between 2011 and 2013 caused analysts to ask whether market prices reflected "true" underlying cost decreases or excessive additions to industry-wide manufacturing capacity. This paper examines this issue by developing a method by which to study cost dynamics in a maturing industry such as the solar PV module sector. Our analysis also yields an estimate of the trajectory of module prices going forward.

The main conceptual contribution of our analysis is the derivation of the *Economically Sustainable Price* (ESP) measure. We demonstrate that the ESP can be interpreted both as the long-run marginal cost of manufacturing one unit of output and as the price that would emerge along a long-run equilibrium trajectory of manufacturing capacity additions. Accordingly, the ESP provides a benchmark against which observed prices could be compared to determine whether the market was in long-run equilibrium at different points in time. In the context of the solar module market, the ESP concept allows us to gauge to what extent the substantial large price reductions observed in recent years were matched by corresponding large cost reductions.

One obstacle in using the ESP concept is that the cost data needed to estimate it are usually unavailable to outside observers. To overcome this, we provide a cost inference procedure to derive the ESP from publicly available financial accounting data. We apply this method to data from solar PV module manufacturers to infer quarterly ESPs and test whether observed ASPs were consistent with a market in long-run equilibrium. While ASPs and ESPs are statistically indistinguishable for most of our sample periods, they are significantly different in at least four quarters in 2012 and 2013, when module prices diverged most sharply from the 80% learning curve.[48] Accordingly, we conclude that the recently observed price reductions reflect a market dynamic driven partly by overcapacity rather than mere cost reductions.

Our assertion that the market was characterized by overcapacity might be interpreted as in support of claims that the Chinese firms that have invested in recent capacity expansions "dumped" their modules in the U.S. and European markets. However, our ESP estimates make this interpretation implausible. Our estimated core manufacturing costs are \$0.71/W and \$0.55/W in 2012 and 2013, respectively, suggesting that modules sold at prices above those benchmarks were unlikely to have been dumped into export markets.[49] Our comment on dumping allegations is corroborated by the minimum selling price of \$0.74/W for Chinese solar panels in the European Union that was set by Chinese and E.U. negotiators in the middle of 2013 (Kanter and Bradsher, 2013).

Our cost inferences also provide us with a panel of data with which to estimate how production costs for modules change as a function of time and experience. Using estimated cost decline parameters, we extrapolate a trajectory of future production costs. This path represents our benchmark of the economic industry fundamentals and is interpreted as a trend-line to which equilibrium prices should converge over time. Upon controlling for plant scale and periods in which polysilicon prices were declining steeply, cost data are consistent with a 74% learning curve for core manufacturing costs. Combined with capacity cost declines that occur yearly, we observe a net 73% cost reduction curve. We thus anticipate an ESP equal to 73% of its previous value with each doubling in cumulative production.

[48] This holds with a 10% significance threshold. With a stricter cutoff at a 5% significance level, we reject the equality of the ASP and ESP in two of the periods, Q1-12 and Q3-12.

[49] The annual ESP figures are weighted averages of quarterly core manufacturing costs in each year. For a given year, the measure weights each quarter's observations by the share of annual shipments observed in that quarter.

Though this rate of cost declines is faster than that embedded by the 80% price-based learning curve, our point estimates of future production costs suggest that module prices are unlikely to decrease to the $0.50/W level targeted by the U.S. SunShot program until 2025, at least if the industry's firms were to earn zero economic profits. Even if demand were to equal 60GW/year, our point estimate of cost decreases implies that 2020 costs would remain at $0.56/W. While the SunShot target is covered by 95% prediction intervals, our work suggests that the business-as-usual rate of technological progress may not itself deliver the cost targets envisioned by the SunShot program. The realization of these targets could require either technological advancements in module manufacturing or greater-than-expected reductions in downstream costs, which are collectively termed the “balance-of-system” (BOS) costs.

A Proof of Finding 1

We verify that the sequence of K_t^* given by:

$$P_t^o(K_t^*) = ESP_t \equiv c_t \cdot \Delta + w_t + f_t,$$

is indeed implementable by a sequence of non-negative investments I_t^* if

$$K_t = I_{t-T} + I_{t-T+1} + \ldots + I_{t-1},$$

and $K_0 \leq K_1^*$. The non-negativity constraints are met if $K_{t+1}^* \geq K_t^*$ for $t \geq 1$. This follows from the observation:

$$P_{t+1}^o(K_{t+1}^*) = ESP_{t+1} < ESP_t = P_t^o(K_t^*),$$

combined with the NEC condition requiring that $P_{t+1}^o(K) > P_t^o(K)$ for all K.

It remains to verify that, given the aggregate capacity levels $\{K_t^*\}_{t=1}^{\infty}$, firms will break-even on their investments. Without loss of generality, assume that a particular firm invests in one unit of capacity at time t. The prevailing equilibrium market price in the next T periods is given by $P_{t+\tau}^o(K_{t+\tau}^*) = ESP_{t+\tau}$, with $\tau \in [1, T]$. The firm utilizes this capacity over the next T periods. The pre-tax cash flows of this investment are given by:

$$CFL_t = -v \cdot \eta^t,$$

and for $1 \leq \tau \leq T$,

$$CFL_{t+\tau} = ESP_{t+\tau} - w_{t+\tau} - f_{t+\tau}.$$

Taxable income in period $t + \tau$ becomes:

$$In_{t+\tau} = CFL_{t+\tau} - d_\tau \cdot v \cdot \eta^t.$$

Given a corporate income tax rate of α, the overall NPV of the investment is:

$$NPV = \sum_{\tau=1}^{T} \left[CFL_{t+\tau} - \alpha \cdot I_{t+\tau}\right] \gamma^\tau - \eta^t \cdot v \tag{26}$$

To show that the expression in (26) is indeed zero, we rewrite it as:

$$NPV = (1-\alpha) \cdot \sum_{\tau=1}^{T} \Delta \cdot c_{t+\tau} \cdot \gamma^{\tau} + \alpha \cdot \sum_{\tau=1}^{T} d_{\tau} \cdot \eta^{t} \cdot v \cdot \gamma^{\tau} - \eta^{t} \cdot v, \tag{27}$$

where the tax-factor is as defined in the main text:

$$\Delta = \frac{1 - \alpha \cdot \sum_{\tau=1}^{T} d_{\tau} \cdot \gamma^{\tau}}{1-\alpha}. \tag{28}$$

The second term on the right-hand side of (27) denotes the tax shield. Dividing (27) by $(1-\alpha)$, we obtain:

$$\frac{1}{(1-\alpha)} \cdot NPV = \Delta \left[\sum_{\tau=1}^{T} c_{t+\tau} \cdot \gamma^{\tau} - \eta^{t} \cdot v \right] = 0,$$

because

$$\sum_{\tau=1}^{T} c_{t+\tau} \cdot \gamma^{\tau} = \eta^{t} \cdot \sum_{\tau=1}^{T} c_{\tau} \cdot \gamma^{\tau} = \eta^{t} \cdot v.$$

■

B Data and adjustments for cost inferences

Firms in the study set

Firm	Ticker	2012 Capacity %	2012 Production %
Yingli Green Energy	NYSE: YGE	3.7	6.4
Trina Solar	NYSE: TSL	3.7	4.6
Suntech Power	NYSE: STP	4.0	5.9
Canadian Solar	NASDAQ: CSIQ	3.4	4.8
LDK Solar	NYSE: LDK	2.7	1.5
Hanwha Solar One	NASDAQ: HSOL	2.4	2.6
JA Solar	NASDAQ: JASO	3.3	3.6
ReneSola	NYSE: SOL	1.6	1.7
JinkoSolar	NYSE: JKS	1.9	2.7
China Sunergy	NASDAQ: CSUN	1.3	1.1

Table 6: *Firms included in the study set, including stock tickers, capacity, and production market share in 2012, as recorded by Lux Research (2013). We show 2012 shares since we do not observe the 2013 shares for Suntech Power and LDK Solar.*

Price data

Since prominent sources of module price data (i.e., Bloomberg New Energy Finance and pvXchange) began collecting price data no earlier than Q3-09, the average sales price (ASP) measure we use in our graphs is a composite of several indexes. The measure equally weighs our estimates of in-sample ASPs and a composite of price indexes that we obtain either from Swanson (2011) or the Bloomberg terminal system. For each firm and quarter, we derive firm-specific ASPs as the quotient of revenues and the sales volume (i.e., the module-equivalent shipment level; see Section 3.1.1):

$$ASP_{it}(firm-specific) = \frac{Revenue_{it}}{s_{it}^{ME}} \tag{29}$$

This implies a quarter-specific average in-sample ASP:

$$ASP_t(in-sample) = \sum_i w_{it} ASP_{it}, \tag{30}$$

where the weights in the above summation are in proportion to the firm's share of module-equivalent shipments across the firms in our sample in that quarter.

Our composite of indexes reflects the data available for a particular period. Prior to Q1-10, we use price data included in Swanson (2011). After Q4-10, we use a Bloomberg New Energy Finance index for multi-crystalline silicon module prices. To bridge the gap between the data from Swanson and that available from BNEF, we use the pvXchange Crystalline Modules China Price available from the Bloomberg terminal. We chose this index for Q1-10 through Q3-10 among those from PVXchange and PVinsights and accessible from the Bloomberg terminal system because it offered the best match with the BNEF multi-crystalline silicon module price index over the time periods in which we could observe both indexes. The ASPs on our graphs equal the simple average of our composite index and in-sample ASP measures and reflect an adjustment to 2013 dollars.

In contrast, our tests about whether the market is in long-run equilibrium in a given quarter use only the firm-specific ASP estimates. By using the firm-specific ASP estimates, we can compare a distribution of estimated ASPs and inferred ESPs across firms. Price indexes, on the other hand, provide us with only one quarterly price observation without any information about the spread of ASPs observed. We observe a relatively close match between our in-sample ASPs and the index price data; across the 24 quarters in our sample, we observe a median difference of 12% between the two numbers.

Adjusting facility capacity costs for physical efficiency gains

As the efficiency of solar cells increases, the same physical area of output contains a greater Watt capacity and therefore the capacity cost per Watt decreases. Recalling that $\eta_f = 1$, we modify c_{ft} from its form in (16) to:

$$c_{ft} = \frac{v_f \cdot \frac{eff_{ref}}{eff_t}}{\sum_{\tau=1}^{T} \gamma^\tau}.$$

Here, eff_t and eff_{ref} refer to average efficiency levels in the current and baseline periods, respectively. We use data from Fraunhofer (2012) and reported in Table 7 to adjust our capacity cost estimates.

Deriving module-equivalent capacity

In Section 3.1.2, we discuss two adjustments to the measure of firm-specific equipment capacity costs. One modification adjusts the level of capacity additions by firms. The naive approach implied by (17) assumes that a firm's capital expenditures on equipment are solely applied to the expansion of module manufacturing capacity. However, as we

2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013
12.0%	12.5%	12.5%	12.7%	13.0%	13.1%	13.1%	13.4%	14.5%	14.7%	15.5%

Table 7: *Average crystalline silicon module efficiency, from Fraunhofer (2012). The 2013 figure is an estimate, given reports by Fraunhofer that 2014 average efficiency levels had reached 16%.*

discuss in Section 3.1.2, firms could have expanded capacity for any of the four steps of module manufacturing or any combination thereof. In practice, firms have tended to invest in capacity either for only one of the four steps or for combinations of the four steps that are contiguous to each other and include all upstream steps. The latter observation implies that firms have invested in, for example, cell, wafer, and ingot capacity but not in only cell and ingot capacity. This practical reality implies that there are ten types of what we term *contiguous capacity investment bundles.*[50]

Of course, it is not equally costly to expand each of the ten capacity bundles. To account for these differences, we use a module-equivalent (ME) level of capacity, K^{ME}:

$$K^{ME} = \Sigma_{l=1}^{l=10} K_j \cdot \chi_j \tag{31}$$

Here, the index j refers to the ten contiguous capacity investment bundles. We use quarterly firm-level capacity data from Lux Research (2014) across all steps of the value chain to derive K_j and K^{ME}.[51] We determine the expansion of a particular bundle by (1) taking the minimum of the capacity expansions for all constituent value chain steps and (2) subtracting the expansions recorded for more inclusive bundles. As an example, when calculating the capacity expansion in the "cells and modules" bundle, we know that this increase cannot

[50]The ten bundles are investments in (1) ingots only, (2) wafers only, (3) cells only, (4) modules only, (5) wafers and ingots, (6) cells, wafers, and ingots, (7) cells and wafers, (8) modules, cells, wafers, and ingots, (9) modules, cells, wafers, and (10) modules and cells.

[51]We make several amendments to the records in this dataset based on our findings from firms' press releases and industry analysts. For Yingli, we use a 2014 module capacity of 2.45GW because we exclude tolling facilities to which the firm had access. For JA Solar, we added 150MW of module capacity in 2014 that stem from new capacity in South Africa but not included in the Lux Research dataset. For China Sunergy, we use an 1.155GW module capacity figure reported by Bloomberg New Energy Finance that matches with numbers released by the firm itself. Finally, we use 800MW for Hanwha SolarOne's ingot capacity; while Lux Research (2014) does not include any ingot manufacturing capacity for this firm, previous versions of the same dataset record an ingot capacity of 800MW per year. In addition, the firm reports that it operates ingot manufacturing capacity.

exceed the observed expansion of either cell or module capacity (i.e., the constituent value chain steps). We thus calculate the minimum capacity expansion level observed across these two steps. To avoid double counting capacity expansions in cells and modules, we subtract the capacity expansion observed across the two more inclusive bundles, namely "modules, cells, wafers, and ingots" and "modules, cells, and wafers."

χ_j is an adjustment factor that "marks down" the capacity additions for bundles that do not include all four components of the value chain. We define χ_j as the ratio of the capacity cost for bundle j to the capacity cost for the integrated module capacity investment. We estimate χ_j as the average of the ratios from 2009 to 2016 implied by Table 2.

C Learning curve estimation robustness checks

Accounting for physical efficiency gains

Since a time trend would have accounted for core manufacturing cost reductions due to improved quality, as measured by physical efficiency, we repeat Specifications 1 through 4 with data expressed on a dollar per square meter basis. Since $efficiency = \frac{power}{m^2}$, we convert manufacturing costs from dollars per watt to dollars per square meter by multiplying the former by efficiency. We use average module efficiency levels from Table 7 and change the scale and cumulative output measures to a square meter basis. Table 8 summarizes our estimates; each specification corresponds to the similarly numbered one in Table 4.

	Dependent variable: Log Manufacturing Cost/m^2			
Specification	**1SM**	**2SM**	**3SM**	**4SM**
Intercept	-0.604***	-0.632***	-0.689***	-1.273***
	(0.155)	(0.155)	(0.183)	(0.146)
Cumulative Production (b)	-0.391***	-0.367***	-0.395***	-0.188**
	(0.044)	(0.053)	(0.048)	(0.053)
Firm Scale (b_s)	–	-0.000	-0.000	0.000
		(0.000)	(0.000)	(0.000)
Dummy, PS Regime 4	–	–	–	-0.431***
				(0.106)
Learning Curve Slope (S)	76.3%	77.5%	76.1%	87.8%
	(15.7%)	(20.0%)	(16.7%)	(41.2%)
Adjusted R^2	0.7273	0.7297	0.7337	0.7784
N	213	213	125	125

Table 8: *Estimated coefficients on a constant elasticity learning curve. Entries in parentheses are Driscoll-Kraay standard errors. The intercept should be interpreted as the average of the logarithm of the Q1-08 core manufacturing cost across firms.*
Key to statistical significance: ***: ≤ 0.001; **: ≤ 0.01; *: ≤ 0.05.

Comparing Tables 4 and 8, we observe that the implied learning curve slopes are essentially the same across the two tables, with the slopes on a \$/Watt basis roughly up to 1 – 2% steeper than those on a \$/m^2 basis. The estimates suggest that our results are robust to the specification of output on an efficiency-adjusted basis.

Exclusion of firm-quarter pairs

Since some firms had a small share of modules in their output mix over some of the periods in our panel, we include four robustness checks in which we exclude observations for these firms from such periods. In the first, we drop data from CSUN between Q1-08 and Q3-10. The second drops data from JASO between Q1-08 and Q4-11, while the third drops data from SOL between Q1-08 and Q3-10. Finally, we drop all three sets of observations. We do not list the estimates derived upon dropping these observations, but we do not find any material differences between the learning curve slopes estimated from the full sample and those obtained when using the restricted samples. Though the standard errors change, they do not change in a systematic direction with these exclusions.

References

Afshartous, D., and R. A. Preston (2010), "Confidence intervals for dependent data: equating non-overlap with statistical significance," *Computational Statistics & Data Analysis*, 54(10), 2296–2305.

Arrow, K. (1964), "Optimal Capital Policy, The Cost of Capital and Myopic Decision Rules," *Annals of the Institute of Statistical Mathematics*, 1-2, 21 – 30.

Baldenius, T., A. Nezlobin, and I. Vaysman (2015), "Managerial Performance Evaluation and Real Options," *Working Paper, New York University.*

Besanko, D., U. Doraszelski, and Y. Kryukov (2014), "The Economics of predation: What drives pricing when there is learning-by-doing," *The American Economic Review*, 104(3), 868 – 897.

Bloomberg New Energy Finance (2014), "Bloomberg Solar," Accessed via the Bloomberg Terminal service.

Borenstein, S. (2012), "The Private and Public Economics of Renewable Electricity Generation," *Journal of Economic Perspectives*, 26(1), 67 – 92.

Comello, S. D., and S. J. Reichelstein (2015), "The US Investment Tax Credit for Solar Energy: Alternatives to the Anticipated 2017 Step-Down," *GSB Working Paper No. 3272.*

Dick, A. R. (1991), "Learning by Doing and Dumping in the Semiconductor Industry," *Journal of Law and Economics*, 34(1), pp. 133–159.

Driscoll, J. C., and A. C. Kraay (1998), "Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data," *Review of Economics and Statistics*, 80(4), 549–560.

Dutta, S., and S. Reichelstein (2010), "Decentralized Capacity Management and Internal Pricing," *Review of Accounting Studies*, 15(3), 442 – 478.

Fraunhofer (2012), "Photovoltaics Report," Tech. rep., Fraunhofer Institute for Solar Energy Systems.

Goodrich, A., P. Hacke, Q. Wang, B. Sopori, R. Margolis, T. L. James, and M. Woodhouse (2013a), "A wafer-based monocrystalline silicon photovoltaics road map: Utilizing known technology improvement opportunities for further reductions in manufacturing costs," *Solar Energy Materials and Solar Cells*, 114, 110–135.

Goodrich, A. C., D. M. Powell, T. L. James, M. Woodhouse, and T. Buonassisi (2013b), "Assessing the Drivers of Regional Trends in Solar Photovoltaic Manufacturing," *Energy & Environmental Science*, 6, 2811 – 2821.

Greene, W. H. (2003), *Econometric Analysis*, Pearson.

GTM (2012), "GTM Research PV Technology, Production and Cost Outlook: 2012 – 2016," Tech. rep.

Hoechle, D. (2007), "Robust Standard Errors for Panel Regressions with Cross-Sectional Dependence," *Stata Journal*, 7(3), 281.

Kanter, J., and K. Bradsher (2013), "Europe and China Agree to Settle Solar Panel Fight," New York Times.

Lazard, P. (2009), "Levelized Cost of Energy Analysis," Tech. rep., Lazar Capital Markets Report, version 3.

Lieberman, M. B. (1984), "The Learning Curve and Pricing in the Chemical Processing Industries," *The RAND Journal of Economics*, 15(2), 213 – 228.

Lux Research (2012a), "Q4 2012 Supply Tracker," Available by subscription to Lux Research.

Lux Research (2012b), "Turning Lemons into Lemonade: Opportunities in the Turbulent PV Equipment Market," Available by subscription to Lux Research.

Lux Research (2013), "Q1 2013 Supply Tracker," Available by subscription to Lux Research.

Lux Research (2014), "Q1 2014 Supply Tracker," Available by subscription to Lux Research.

Nezlobin, A. (2012), "Accrual Accounting, Informational Sufficiency, and Equity Valuation," *Journal of Accounting Research*.

Nezlobin, A., M. V. Rajan, and S. Reichelstein (2012), "Dynamics of Rate-of-Return Regulation," *Management Science*, 58(5), 980–995.

Pillai, U., and J. McLaughlin (2013), "A model of competition in the solar panel industry," *Energy Economics*, 40, 32–39.

Powell, D., M. Winkler, H. Choi, C. Simmons, D. Needleman, and T. Buonassisi (2012), "Crystalline silicon photovoltaics: a cost analysis framework for determining technology pathways to reach baseload electricity costs," *Energy & Environmental Science*, 5, 5874 – 5883.

Powell, D., M. Winkler, A. Goodrich, and T. Buonassisi (2013), "Modeling the Cost and Minimum Sustainable Price of Crystalline Silicon Photovoltaic Manufacturing in the United States," *IEEE Journal of Photovoltaics*.

PWC (2012), "Financial Engineering as a Means to Support Jawaharlal Nehru Solar Mission," Tech. rep., PWC.

Rajan, M., and S. Reichelstein (2009), "Depreciation Rules and The Relation Between Marginal and Historical Cost," *Journal of Accounting Research*, 47, 1 – 43.

Reichelstein, S., and M. Yorston (2013), "The Prospects for Cost Competitive Solar PV Power," *Energy Policy*, 55, 117 – 127.

Rogerson, W. (2008), "Inter-Temporal Cost Allocation and Investment Decisions," *Journal of Political Economy*, 116, 931 – 950.

Rogerson, W. (2011), "On the relationship between historic cost, forward looking cost and long run marginal cost," *Review of Network Economics*, 10(2).

Spence, A. M. (1981), "The learning curve and competition," *The Bell Journal of Economics*, 49–70.

Stobaugh, R. B., and P. L. Townsend (1975), "Price forecasting and strategic planning: the case of petrochemicals," *Journal of Marketing Research*, 19–29.

Swanson, R. (2011), "The Silicon Photovoltaic Roadmap," The Stanford Energy Seminar.

Timilsina, G., L. Kurdgelashvili, and P. Narbel (2011), "A Review of Solar Energy: Markets, Economics and Policies," Tech. rep., The World Bank Development Research Group, Environment and Energy Team.